

Security Level:

# Powering the Cloud — Introduction to CloudEngine Series Switches

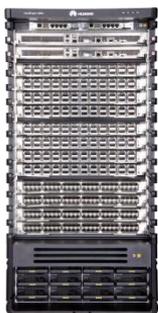
[www.huawei.com](http://www.huawei.com)

HUAWEI TECHNOLOGIES CO., LTD.



# Huawei CloudEngine series switches — high-performance cloud switches designed for next-generation data centers and high-end campus networks — support a rich variety of service features.

## Flagship core switches



CE12812



CE12808



CE12804

## High-performance TOR switches



CE6850-48T4Q-EI



CE6850-48S4Q-EI



CE5850-48T4S2Q-EI

- Data center core switch — CloudEngine 12800 series
- Data center TOR switch — CloudEngine 6800/5800 series
- Mature and stable software platform — VRP8

# CloudEngine 12800 Series

— Core switches for data center and high-end campus networks

# Contents

- **Next-Generation Core Switch—CloudEngine 12800**
  - Next-Generation Core Engine to Provide the Highest Performance in the World
  - Advanced Architecture to Enhance Efficiency and Reliability
  - Comprehensive Virtualization Capabilities to Support Flexible Deployment
  - Innovative Energy Saving Technologies to Build a Green Data Center

# The Highest Performance in the World



**8\*100 GE** line card

**24\*40 GE/96\*10 GE**  
(breakout) line card

	CloudEngine 12800	Average
Line card capability	1 Tbps	2 times 480 Gbps
Switching capacity	48 Tbps	3 times 18 Tbps
Slot bandwidth	2 Tbps	4 times 500 Gbps

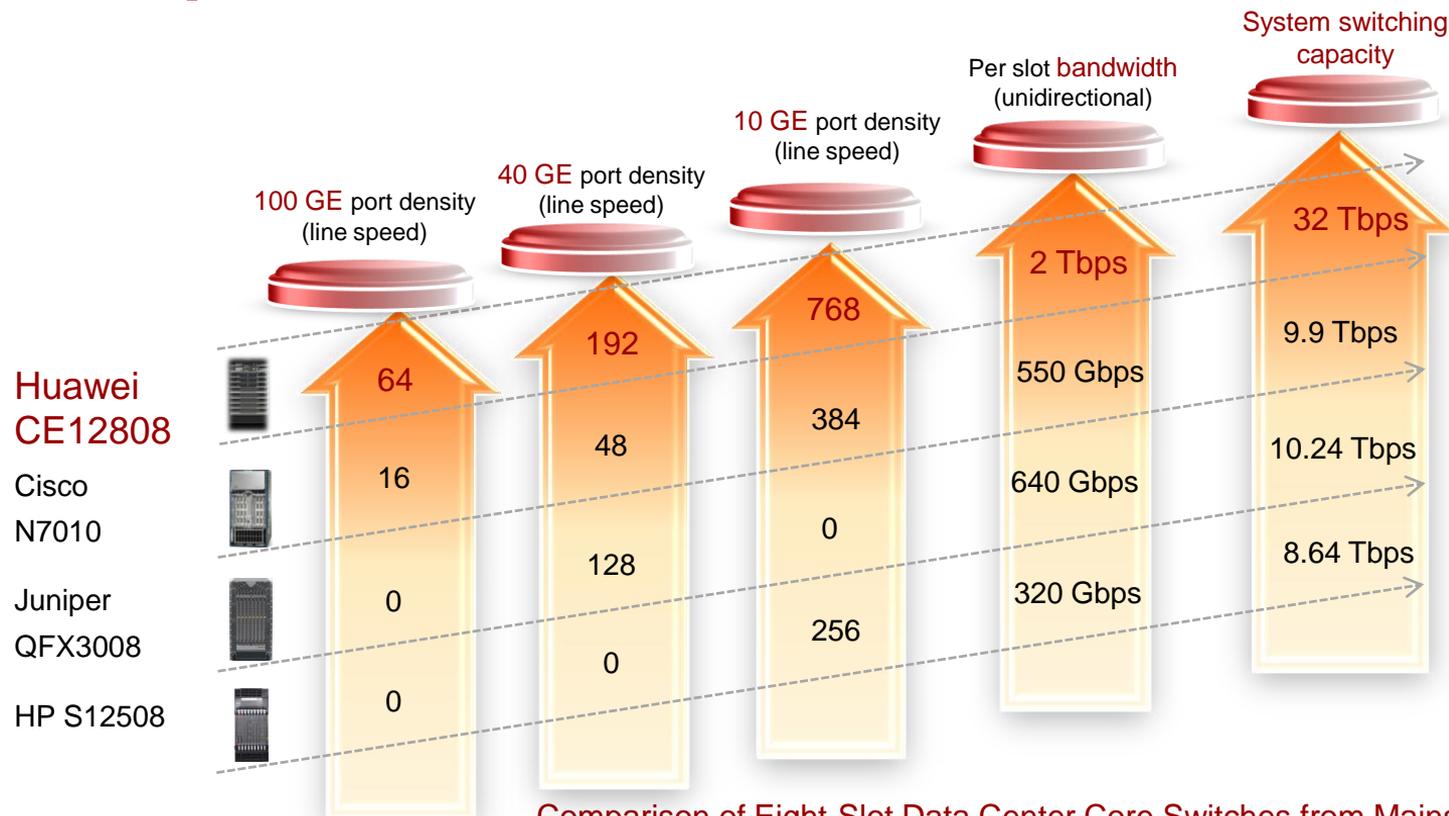
24\*40 GE/96\*10 GE line card, forwarding at line-speed

Multi-level switch fabric, supports large capacity and future expansion

Orthogonal architecture, provides high speed and can be upgraded to 4Tbps in the future

Industry-leading scalable fabric, a stable network platform sustaining **10 years** of evolution

# Comparison with Mainstream Vendors



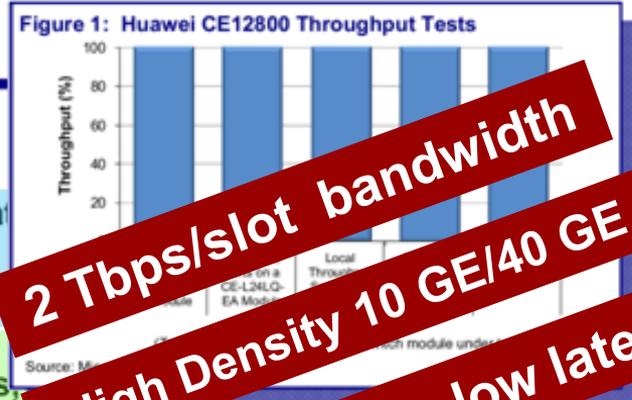
Comparison of Eight-Slot Data Center Core Switches from Mainstream Vendors

# CE12800 Test Report from Miercom



## Key findings and conclusions:

- Huawei CloudEngine 12812 proved full line rate throughput at 2 Tbps per line card without loss using traffic pairs on the same line card as well as traversing cards
- CE12800 is a high performance, scalable core switch supporting up to 1152\*10GbE or 288\*40GbE port connections, up to 47.52 Tbps of line rate traffic and 28,800 Mpps forwarding rate
- Ultra low latency of 2-5 usec and low jitter even during high volume traffic tests and regardless of frame size



**2 Tbps/slot bandwidth**

**High Density 10 GE/40 GE**

**2-5 usec — Ultra-low latency**



# World's Highest Switching Capacity: The CE12800 Series

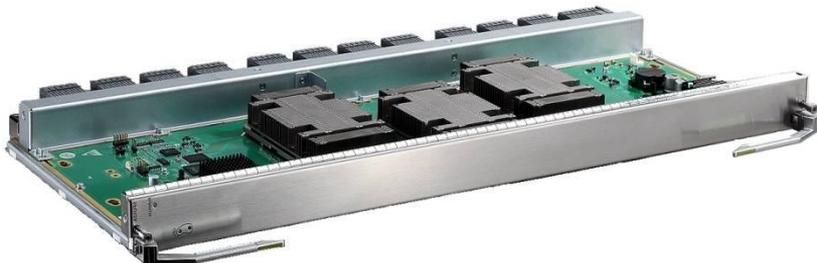
Item	CE12804	CE12808	CE12812	Industry Average
Switching capacity	16 T	32 T	48 T	18.7 T
Number of line card slots	4	8	12	16
Per slot bandwidth	2 T (unidirectional), scalable to 4 T in the future (unidirectional)			550 G
Card Line type	24*40 GE/96*10 GE, 48/24*10 GE, 8*100 GE (2013Q2), all are line-speed cards.			48*10 GE
System power configuration	5400 W (2 + 2 backup)	10800 W (4 + 4 backup)	16200 W (6 +6 backup)	15000 W (4 +4 backup)
Dimensions (W x D x H)	442 x 813 x 486 442 x 969 x 486 (including cabling rack and ejector lever)	442 x 813 x 753 442 x 969 x 753 (including cabling rack and ejector lever)	442 x 813 x 975 442 x 969 x 975 (including cabling rack and ejector lever)	439 x 841 x 1105 (25RU)
System power supply	110/220 VAC			110/220VAC
Port density	192*10 GE 96*40 GE/384*10 GE 32*100 GE	384*10 GE 192*40 GE/768*10 GE 64*100 GE	576*10 GE 288*40 GE/1152*10 GE 96*100 GE	768*10 GE 96*40 GE 32*100 GE
Buffer capacity	Average 100 ms/port (10 GE, 40 GE, and 100 GE have super large buffer.)			14 ms/port

**The CE12800 switches are high-performance core switches designed for data center networks and high-end campus networks. The CE12800 has the highest switching capacity in the industry and super high-density line speed cards that can meet network needs in long term.**

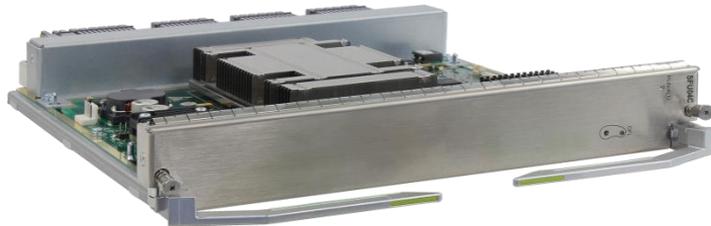
# High-Performance SFU



DE01SFU08C — Independent SFU for CE12808



DE01SFU12C — Independent SFU for CE12812



DE01SFU04C — Independent SFU for CE12804

## SFU

- All CloudEngine 12800 switches have independent SFUs.
- Support a maximum of six SFUs. (N +M backup)
- For high speed 10 GE and 40 GE cards, five SFUs can provide line-speed switching.

	Huawei	CISCO	HP
Switching capacity comparison (unidirectional)	2 T/slot (Scalable to 4T/slot)	550 G/slot	320G/slot

# High-Performance MPU and CMU



CE-MPUA



CE-MUA

## MPUA

- MPUA is the MPU for all CloudEngine 12800 switches.
- Functions:
  - Control and management of system plane
  - Processing of routing protocol packets
  - Control, maintenance, and exchanging of information between SFUs and LPUs cards
  - Management and maintenance according to instructions from users

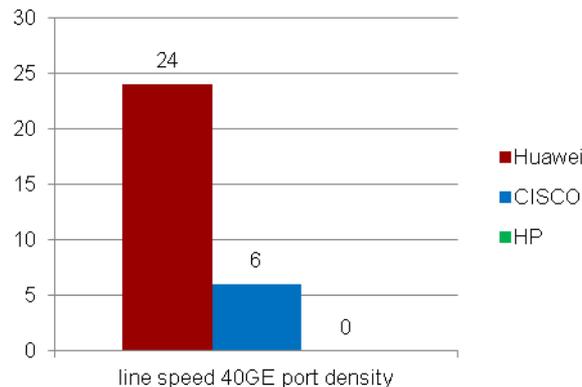
## CMUA

- CMUA is the CMU for all CloudEngine 12800 switches.
- Functions:
  - Collects real-time information about device operation and exchanges data with the management system through high speed channels
  - Provides an independent out-of-band monitoring plane
  - Realizes remote powering on of cards, hardware upgrades, monitoring of temperature, voltage, power, facility management, and fault diagnosis

# High-Density 24\*40 GE/96\*10 G Cards



Orientation	Highlights:	
Provides high-density 10 GE ports in data centers or campus networks	MPLS VPN	12 G large buffer for each card (100 ms per port)
	TRILL: 512 nodes, 32 equal-cost links.	Flex Port: Each 40 GE port can be used as four 10 GE ports.
	MAC: 216 K	Green Port: 3.9 W power consumption for each port (15.7 W for each 40 GE port)



# High-Density 48\*10 GE & 24\*10 GE Cards

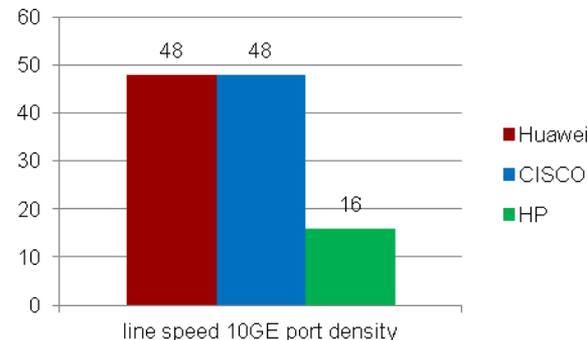


48\*10 GE card



24\*10 GE card

Orientation	Highlights:	
Provides high-density 10 GE ports in data centers or campus networks.	MPLS VPN	6 G/3 G large buffer for each card (100 ms per port)
	TRILL: 512 nodes, 32 equal-cost links.	Flex Port: Each 10 GE port can be used as GE port
	MAC: 216 K	Green Port: 5.5 W power consumption for each 10 GE port



# Flex Port, be Applicable for Different Scenario

## 10GE Flex Port



**SFP-T(RJ45)**  
GE Electrics



**SFP**  
GE Optics  
Connection



**10GBase-CR**  
Copper; Low  
cost connection  
in 10m distance



**SFP+ ,**  
10GE Optics  
Connection



## 40GE Flex Port



**QSFP+ ,**  
150m,10Km  
connection



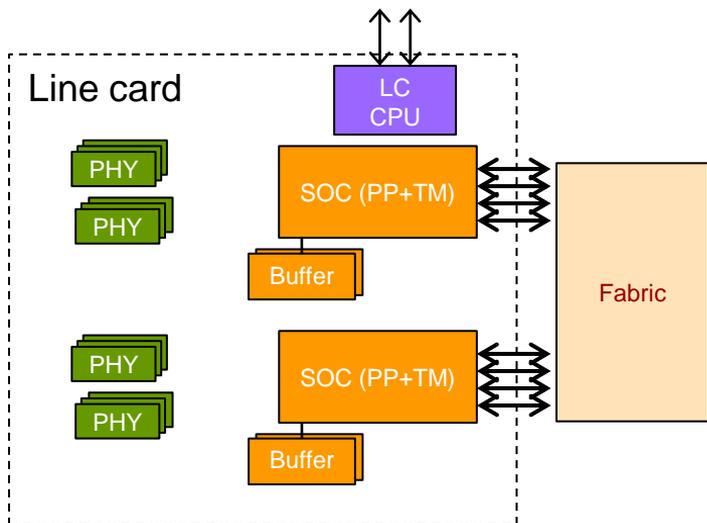
**10GBase-CR4**  
Connection in 5m  
distance



**Optics and copper**  
**1 to 4 breakout,**  
connect with 10GE ports



# High-Performance Line Card: CE12800



High-integration and high-performance line cards have significant advantages in service processing, forwarding, buffering, and power consumption.

## High-performance and high-density ports

- The use of latest generation of ASIC and 45 nm processing technique reduces power consumption.
- 240 Gbps per PP, 1/2/4 ASICs per card
- 48\*10 GE, 24\*40 GE, high density

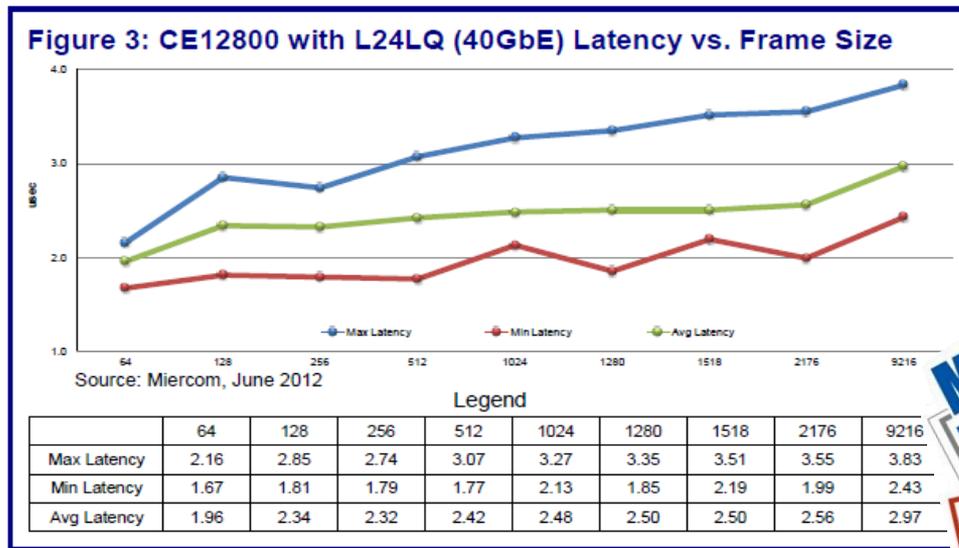
## High integration

- Integration of PP and TM provides unified service processing, buffering, and uplink and downlink switching, greatly reducing packet loss caused by performance bottlenecks.
- A single PP can complete service processing, effectively solving problems of inter-chip service processing.

## Large buffer

- The largest buffer in the industry, 100 ms buffer for 10 GE / 40 GE / 100 GE ports

# 2 $\mu$ s Super Low Forwarding Latency on CE12800



- In store-and-forward mode, the latency on a 40 GE port is **2.45  $\mu$ s**, on a 10 GE port is **3.09 $\mu$ s**.
- The average latency for similar switches from other vendors is about 10  $\mu$ s in the case of 9216-byte packet forwarding. The excellent design of CE12800 reduces the number to 3  $\mu$ s.

# Smooth Evolution to 100 T

## Specifications in 2012

Switching capacity: **48 T**  
Line card capacity: 1 T  
Per slot bandwidth: 2 T

## Same chassis

- Uniform platform
- Uniform system power configuration
- Uniform heat dissipation module

## Specifications in 2015

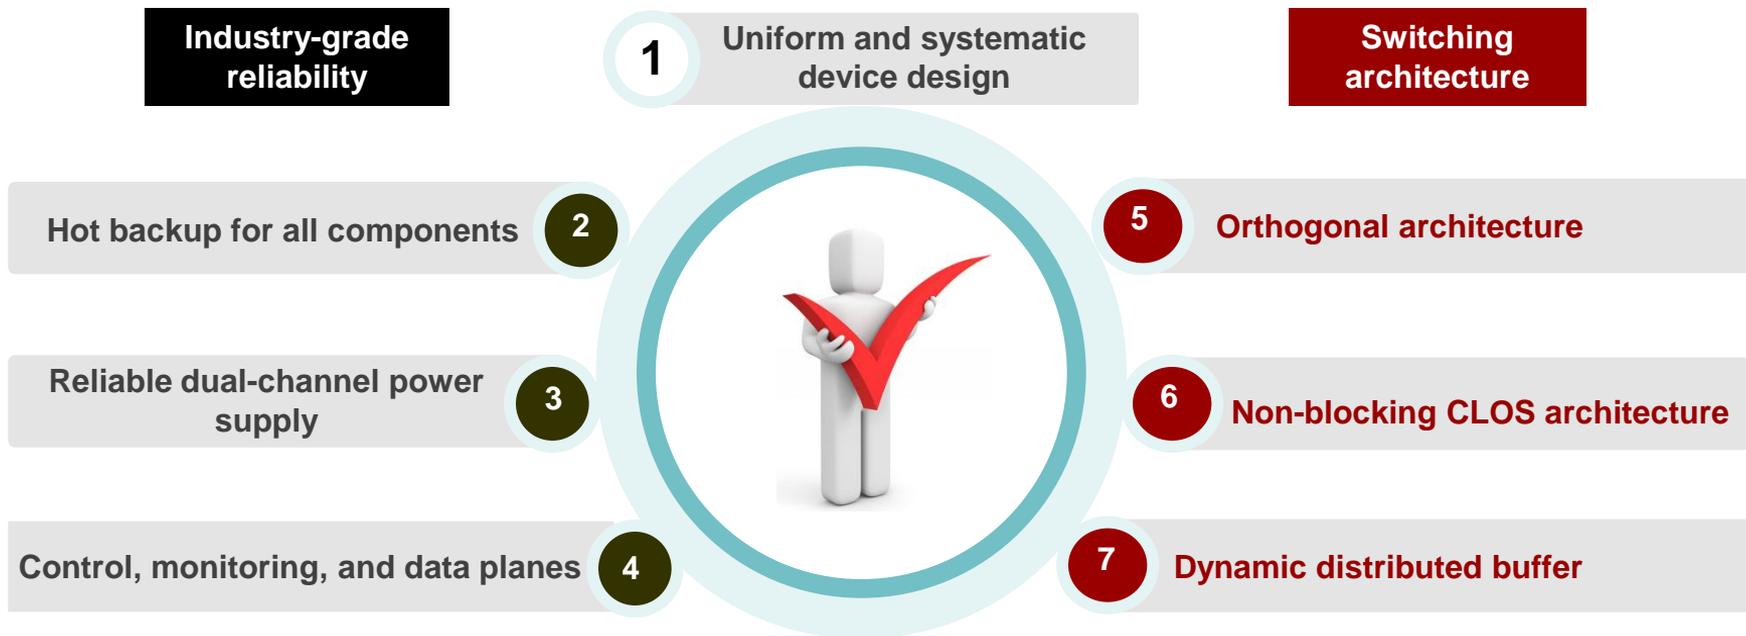
Switching capacity: **115 T**  
Line card capacity: 4 T  
Per slot bandwidth: 4.8 T

- High specifications: 2 T bandwidth per slot and No. 1 10 GE/40 GE port density in the industry meet the development needs of large capacity data centers.
- Smooth evolution: 25 G Serdes was reserved in the system, which helps realize smooth upgrade to 4 T (40 \* 100 GE) card in the future.

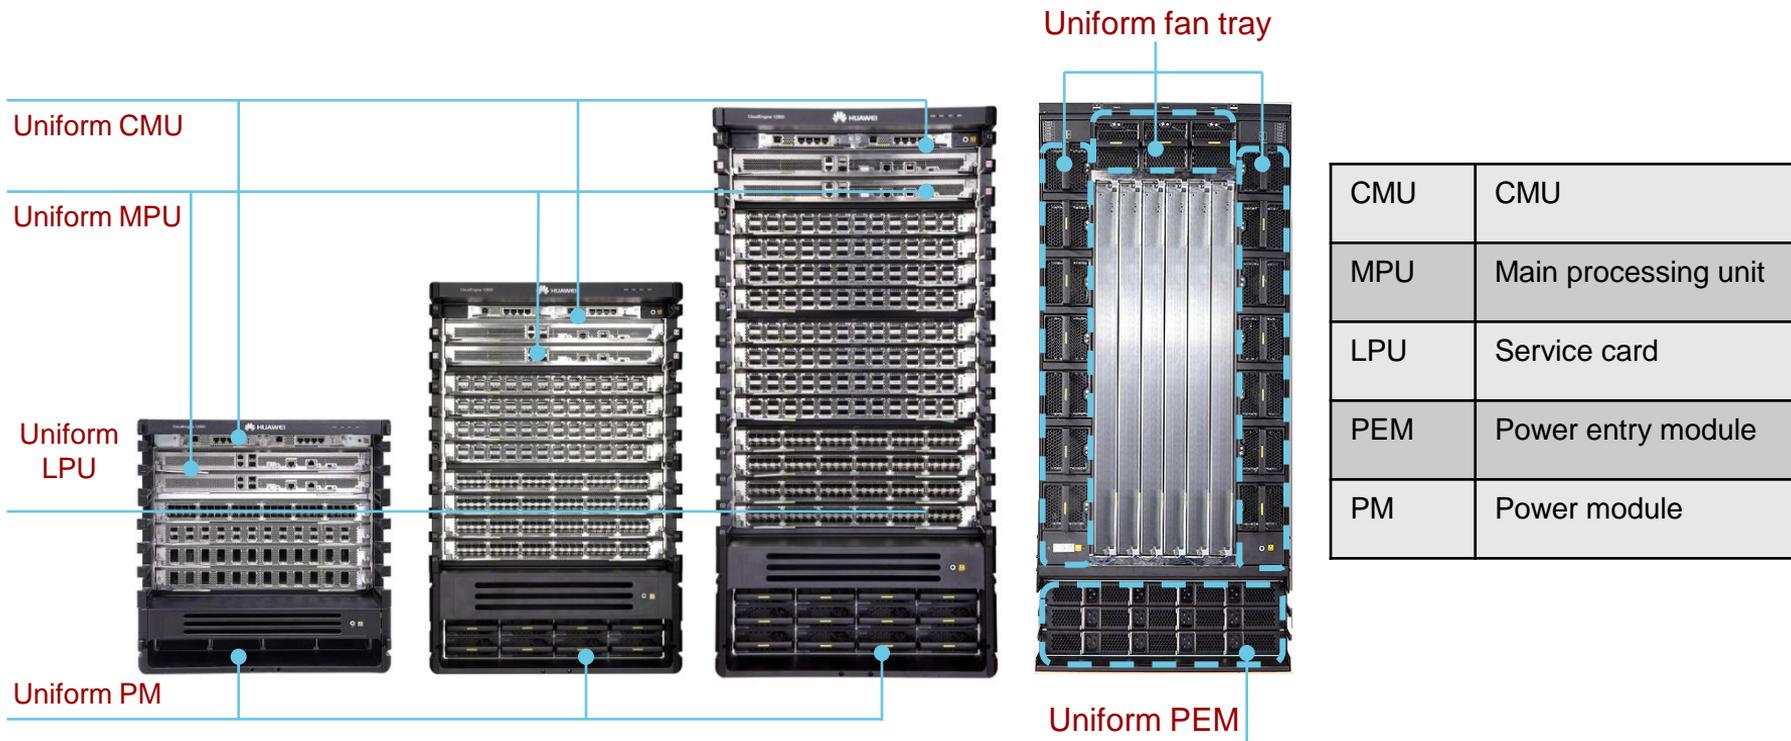
# Contents

- **Next-Generation Core Switch—CloudEngine 12800**
  - Next-Generation Core Engine to Provide the Highest Performance in the World
  - *Advanced Architecture to Enhance Efficiency and Reliability*
  - Comprehensive Virtualization Capabilities to Support Flexible Deployment
  - Innovative Energy Saving Technologies to Build a Green Data Center

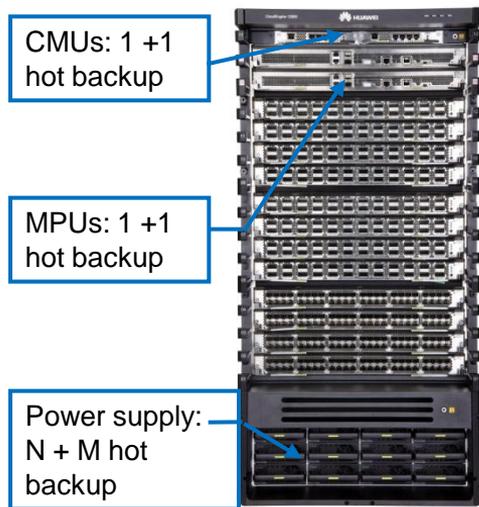
# Advanced Architecture for High Performance



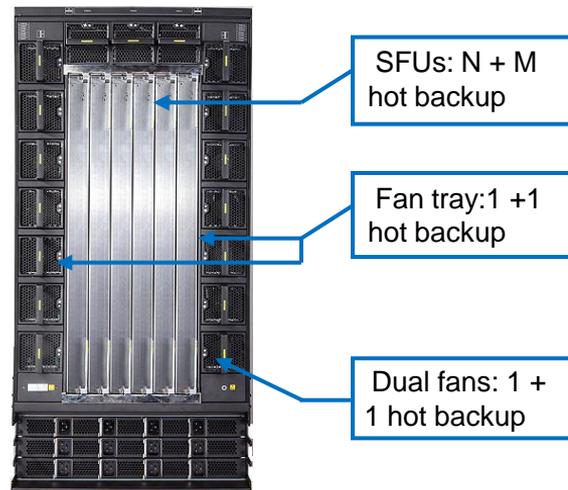
# Uniform and Systematic Device Design



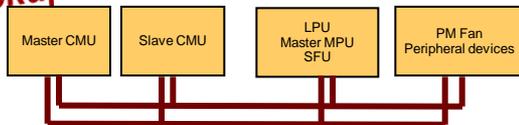
# Industry-Grade Reliability — Hot backup for All



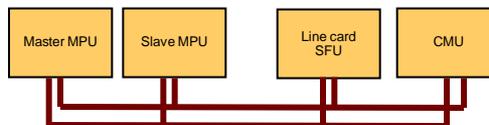
Hot backup of the system



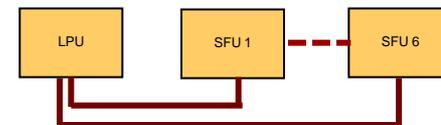
Hot backup in the system



Double CAN MonitorBus

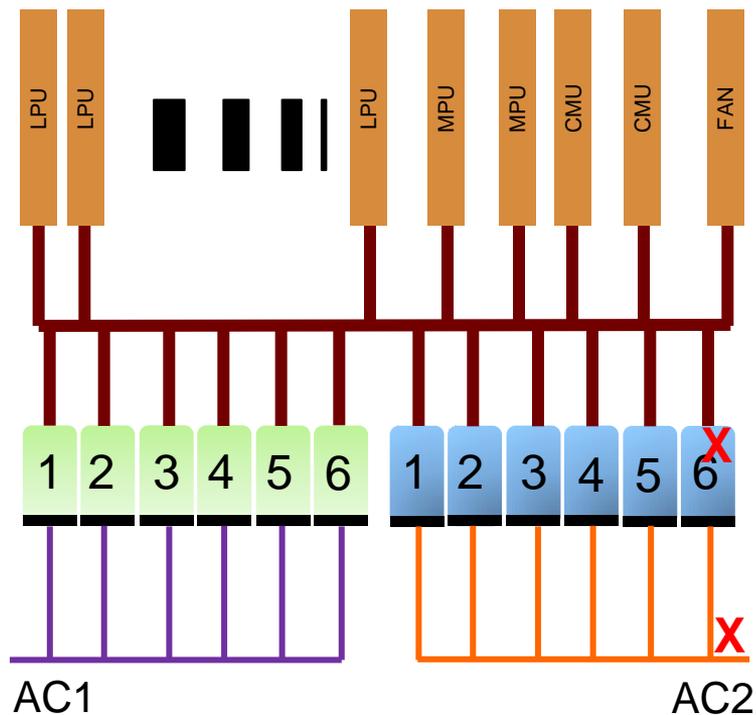


Double GE management bus



Multi-link high-speed data bus

# Industry-Grade Reliability — Dual Power Inputs



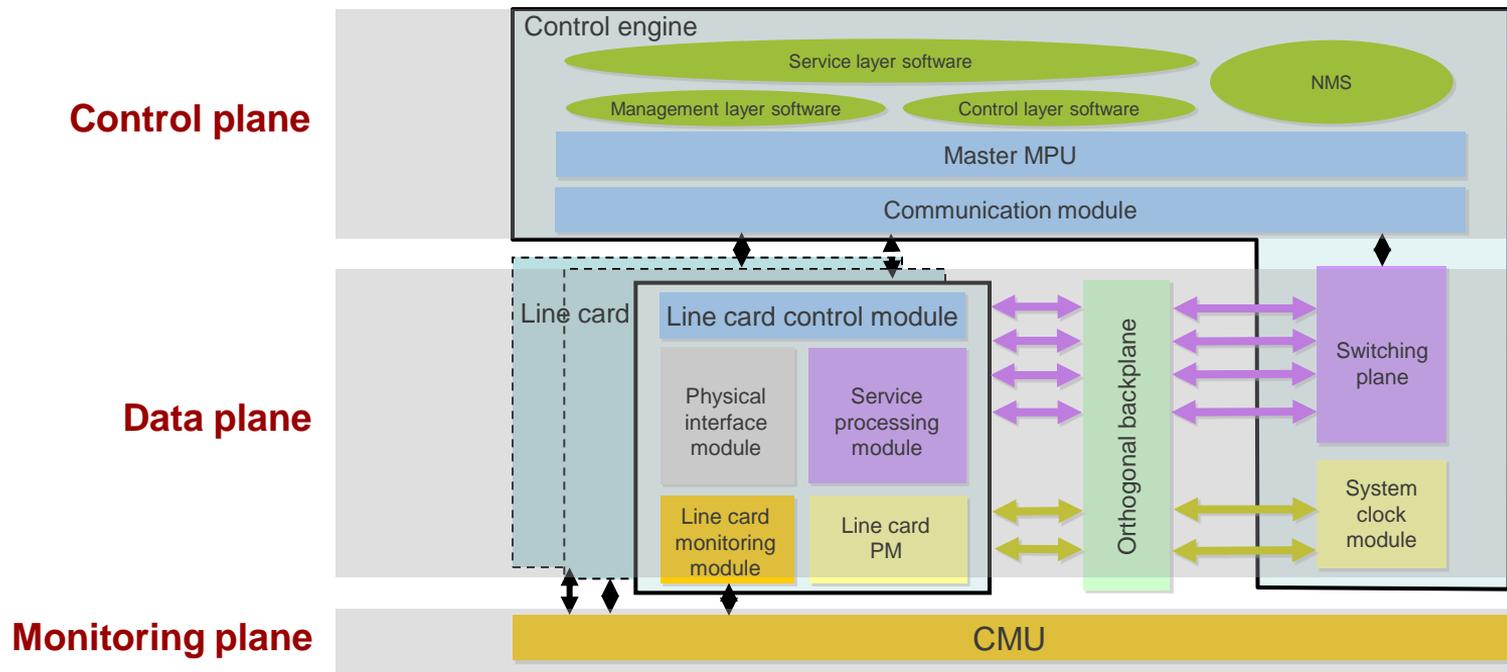
## Efficient and reliable power supply architecture

- Dual power inputs, N + M backup (CE12812/12808/12804 can support a maximum of 12/8/4 PMs respectively)
- Dynamic dormancy improves the power conversion rate by increasing the load rate of non-dormant PMs.

## Highly efficient new - generation PM

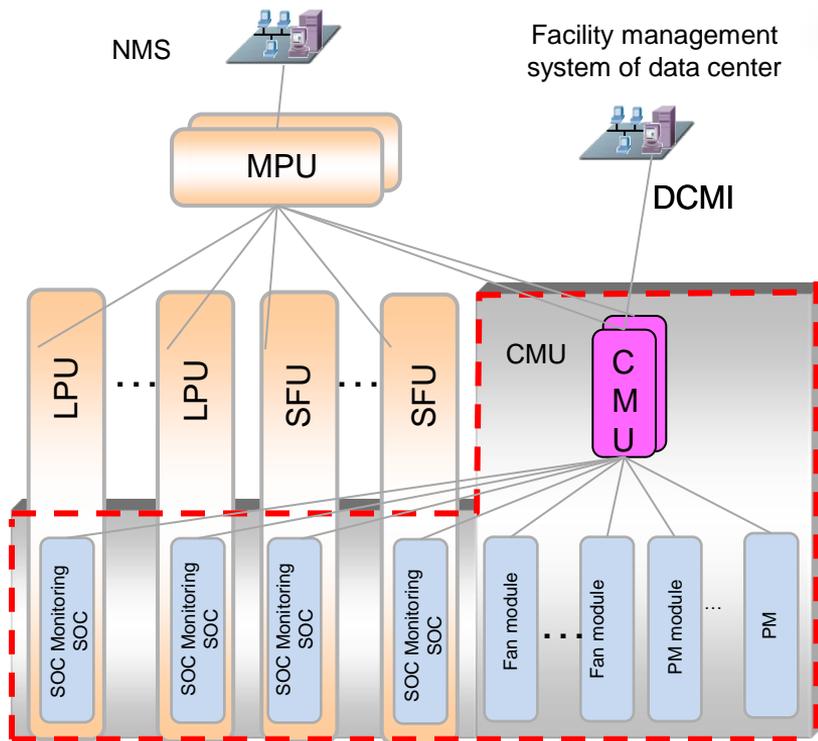
- Power conversion rate  $\geq 96\%$ . Power conversion loss is greatly reduced.
- Wide voltage range: 90 to 290 VAC. Power fluctuation is greatly alleviated.
- Automatic power distribution avoids power loss caused by even power load on some PMs.
- Smart management realizes remote management, dormancy control, and realtime power consumption detection to reduce power consumption.

# Industry-Grade Reliability — Independent Planes



Separate control plane, data plane, and monitoring plane improve system reliability and ensure service continuity.

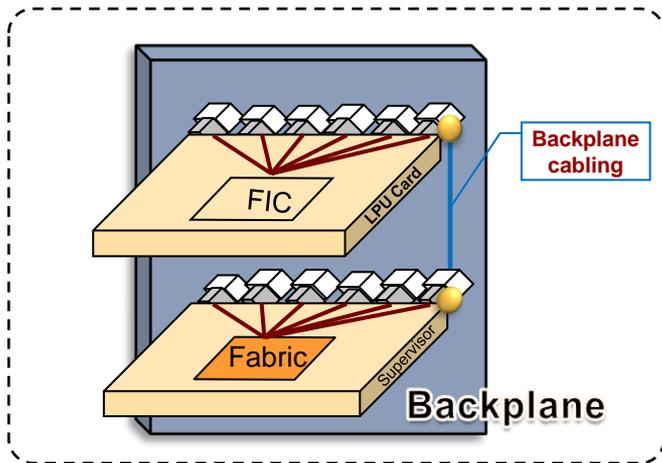
# Independent Out-Band Monitoring Management



## Innovative highly reliable out-band monitoring plane

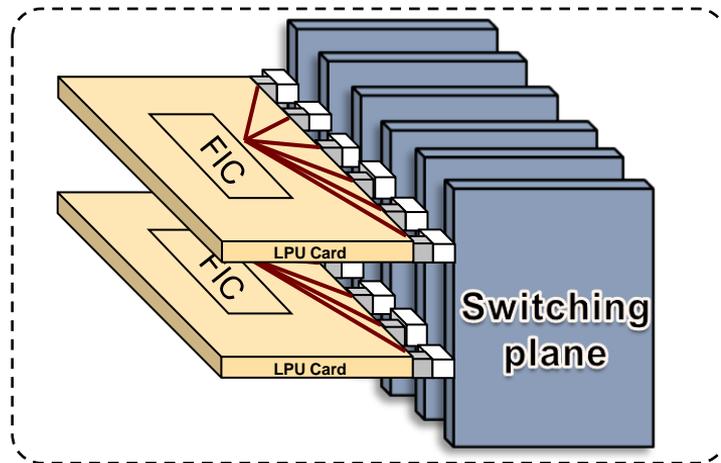
- **Creative monitoring management architecture:** The out-band management and monitoring plane is independent of service and control planes. It can report device status and fault rectification independently, providing **real zero touch** maintenance and management.
- **Highly reliable monitoring architecture:** 1 + 1 backup of CMUs, 1 + 1 backup of double CAN monitor buses.
- **Remote management:** Comprehensive remote monitoring realizes remote card power on and off, device reset, firmware upgrade, temperature/voltage/power monitoring, and fault detection.
- **Standards compliance:** the first switch in the industry that complies with data center DCM1.0 and IPM2.9 management regulations.

# Full Orthogonal Architecture for High Bandwidth, Flexibility and Expansion



Traditional architecture

VS

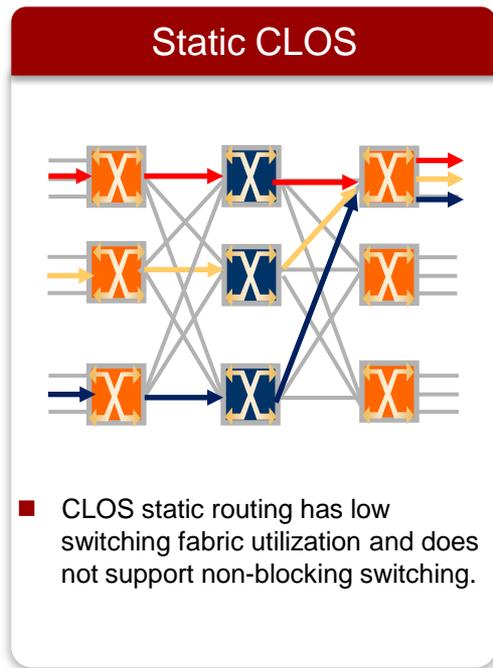


Orthogonal architecture

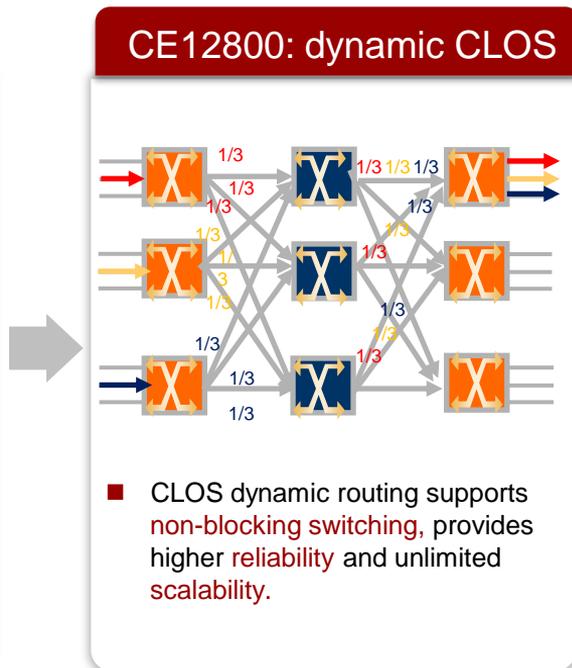
## Excellent full orthogonal architecture

- In core switches, cable connection between line cards and switching fabrics is an important factor affecting slot bandwidth. The longer the backplane cable and its speed, the greater the loss.
- The orthogonal architecture of CE12800 removes the need for backplane cables, which greatly improves the system bandwidth and enhances its evolutionary capability.

# Non-Blocking CLOS Switching Architecture



**Blocking**



**Non-blocking**

- Three-level CLOS non-blocking switching architecture

- Cell switching improves large-packet forwarding efficiency.
- Perfect VoQ mechanism prevents packet loss due to congestion.

- Dynamic routing and distributed arbitration

- Dynamic routing is not affected by traffic model and avoids unbalanced traffic distribution caused by static routing based on hash algorithms.
- Distributed arbitration prevents packets loss during active/standby switchover in central arbitration.

# Advantages of Dynamic CLOS over Static CLOS



*Individual route selection causes congestion, leading to low link utilization.*

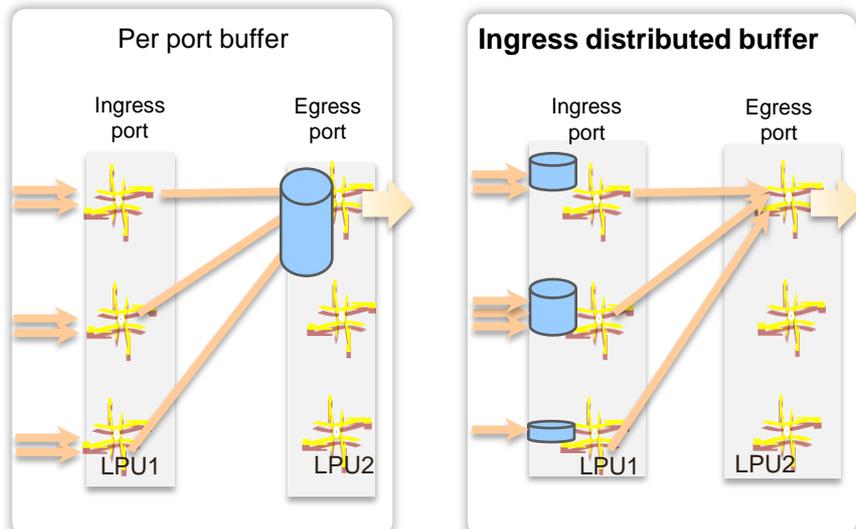


*Unified route scheduling ensures order and high link utilization.*

	Static CLOS Switch Fabric	Dynamic CLOS Switch Fabric
Mechanism	Route selection among HigiG links (four Serdes Trunks) based on hash algorithms, packet switching.	Route selection is based on arbitration, cell switching.
Fairness of packet scheduling	HOL blocking may cause packet loss in the downlink interface card.	VoQ eliminates HOL blocking. Routing based on arbitration avoids packet loss.
Relevance to traffic model	Hash algorithms are based on the source and destination addresses. The hash result is closely related to the traffic model.	Route selection based on arbitration ensures balanced use of links.
Traffic capability	Restricted by HigiG link rate. Maximum rate is 40GE.	The packets are transformed into cells, which are distributed to all Serdes. The traffic rate depends on the rate of a port.
Fault tolerance capability	Based on HigiG, the breakdown of some Serdes in the HigiG may lead to the disconnection of the whole HigiG link. The fault rate is four times higher than for dynamic routing.	Based on Serdes, with a low fault rate.
Link utilization	Route selection based on Hash algorithms may cause unbalanced loads on links: some links are heavily congested whereas others are idle.	Route selection based on arbitration realizes 100% link utilization.
<b>Summary</b>	<b>Applicable only to small-scale and lightly loading networks due to many drawbacks.</b>	<b>High quality switching fabric that can be deployed at data centers and high-end campus network.</b>

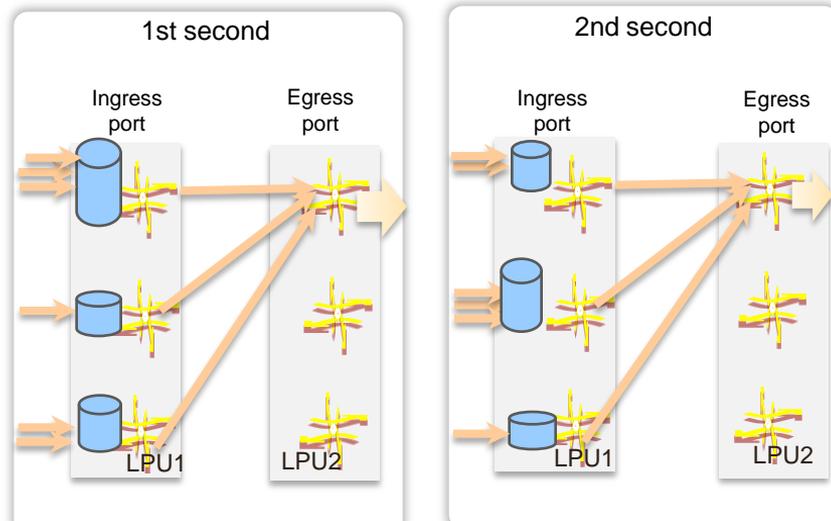
# Dynamic Distributed Buffer for Non-blocking

**Distributed buffer, better absorbs burst traffic**



**Ingress distributed buffer, mitigates buffer load on a single port**

**Dynamic buffer, improves buffer utilization**



**On-demand buffer allocation, dynamic adjustment**



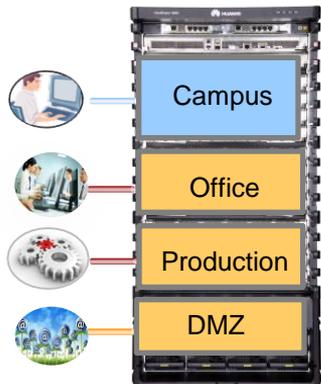
# Contents

- **Next-Generation Core Switch—CloudEngine 12800**
  - Next-Generation Core Engine to Provide the Highest Performance in the World
  - Advanced Architecture to Enhance Efficiency and Reliability
  - Comprehensive Virtualization Capabilities to Support Flexible Deployment
  - Innovative Energy Saving Technologies to Build a Green Data Center

# Virtualized Fabric: VS & CSS Reduce Investment

## Network Virtualization

### VS (Virtual System)

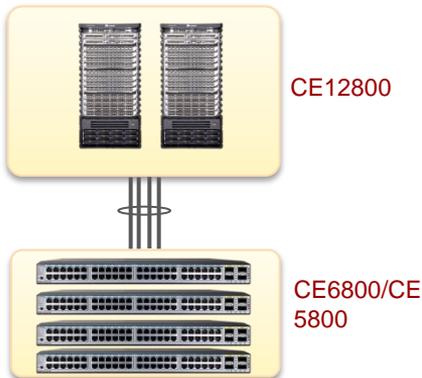


CE12800

#### ■ Network Resource Virtualization:

- VS supports up to **8** virtual switches in one chassis
- CSS supports 2 chassis in one cluster, scalable to **4** (Q2, 2013)

### CSS (Cluster Switch System)



Combine CSS & VS to manage switches more flexibly

## Benefits

### Lower CAPEX

- 
- VS: Fewer devices, higher utilization
  - CSS: 100% bandwidth utilization

### Lower OPEX

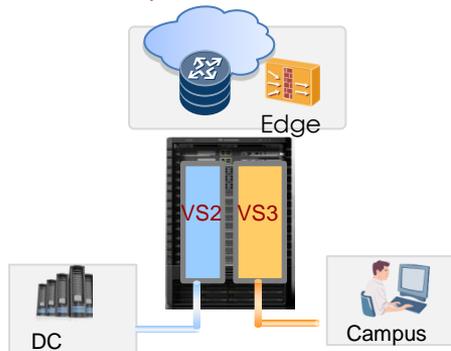
- 
- VS: Fewer physical nodes to maintain
  - CSS: simpler network management

### Secure, Reliable Network

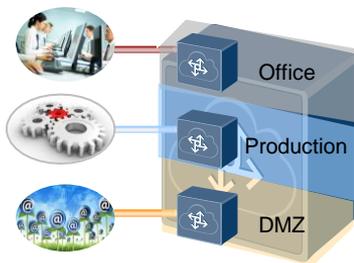
- 
- VS: network and faults isolation
  - CSS: loop-free reliable network

# VS Reduces Required Investment

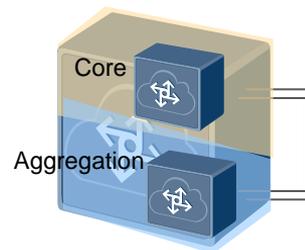
One core device for campus and DC



Multi-zone/multi-tenant



Integrated aggregation and core layer



## Lower CAPEX

- 1:8 virtualization reduces hardware cost.
- On-demand resource allocation improves device utilization.

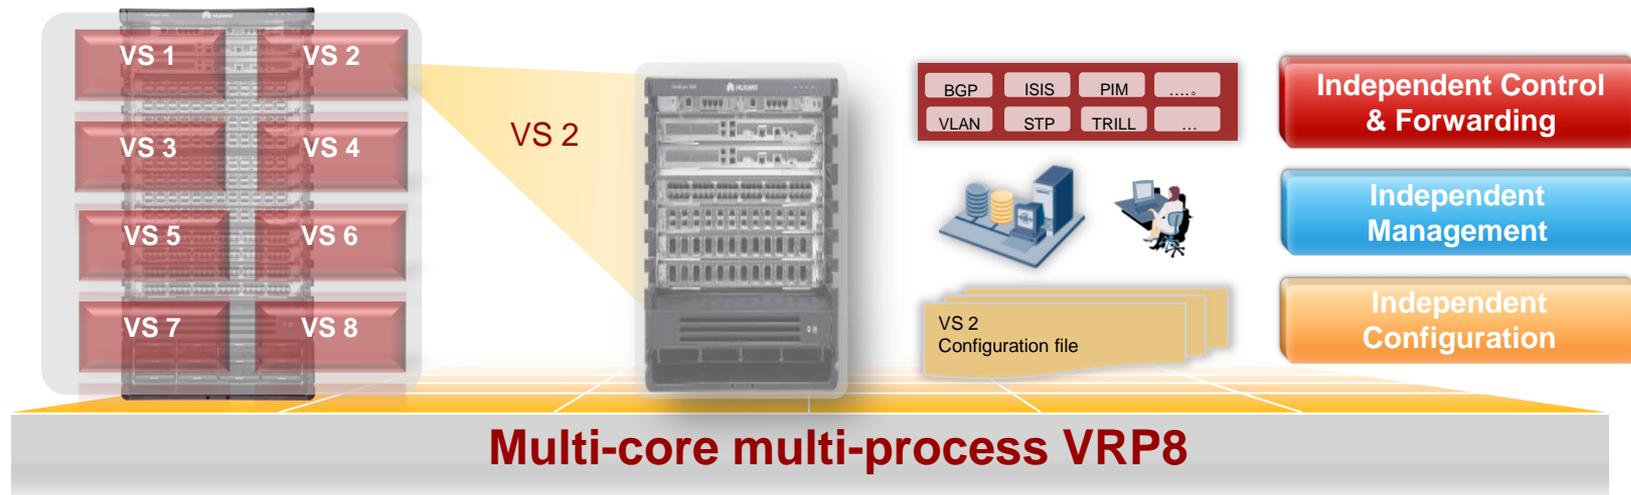
## Lower OPEX

- Multiple zones/tenants share physical devices to reduce power consumption.
- Save physical space.
- Reduce physical nodes to maintain.

## Secure Network

- Network is divided into multiple zones to enhance security.
- Faults on one VS does not affect other VSs.

# VS: Independent Virtual Systems



## VS Components

- Virtual control and service plane
- Virtual forwarding plane
- Virtual management plane

## System Resource Allocation

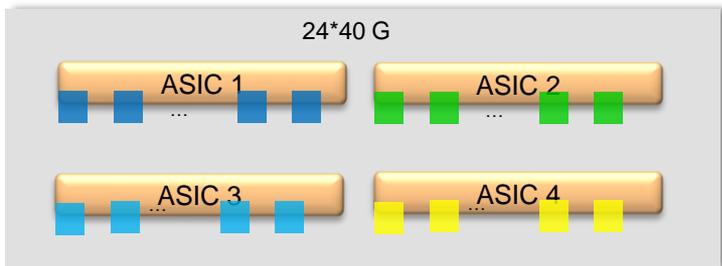
- Per VS physical card and interface allocation
- Per VS system resource specifications

## Independent VS Operation

- Independent configuration, management, and maintenance
- Independent protocol processing
- Isolation from other VSs

# VS: Flexible Resource Allocation

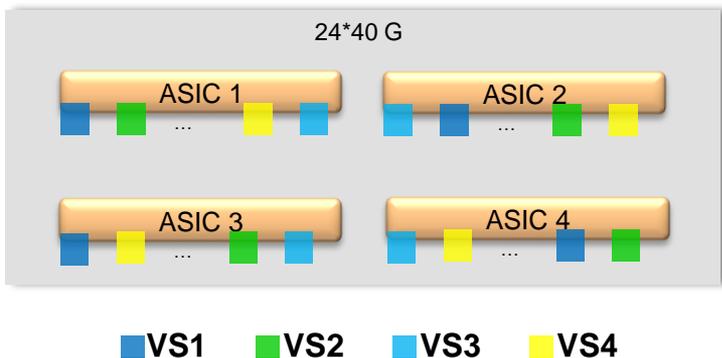
## Per-port group allocation



## ● Per-port group allocation

- A group of ports are allocated to one VS.
- The VS exclusively uses system service specifications.
- A VS can enable all services.

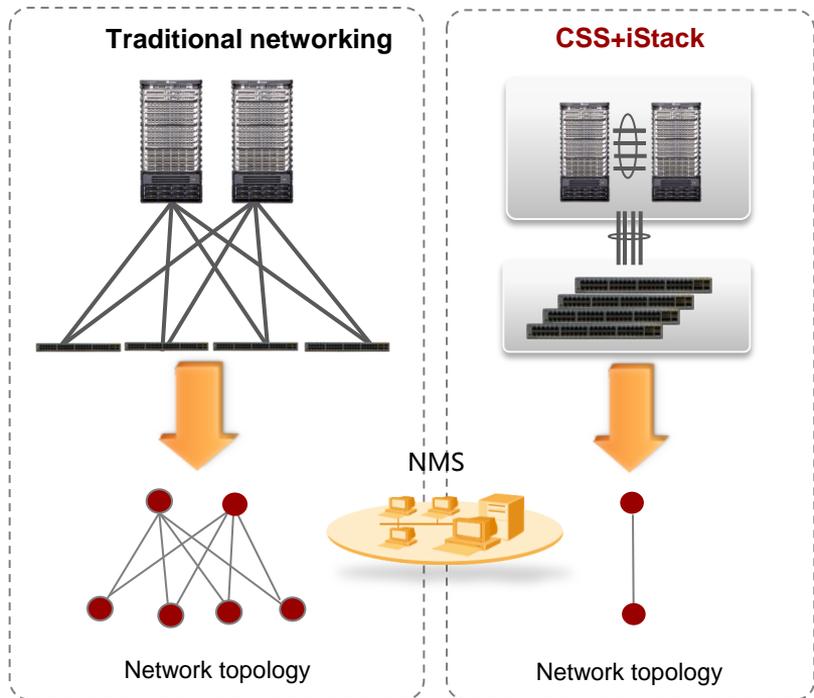
## Per-port allocation



## ● Per-port allocation

- Any port can be allocated to any VS.
- VSes share system service specifications, and some features can only be enabled in one VS.
- Some services (multicast, MPLS, and TRILL) can only be enabled in one VS.

# CSS & iStack Simplify Network Topology & Mgmt.



## Simplify O&M and reduce OPEX

- Fewer network nodes are deployed, simplifying network management.
- Master switch synchronizes configuration file to other member switches, simplifying device operation and configuration.

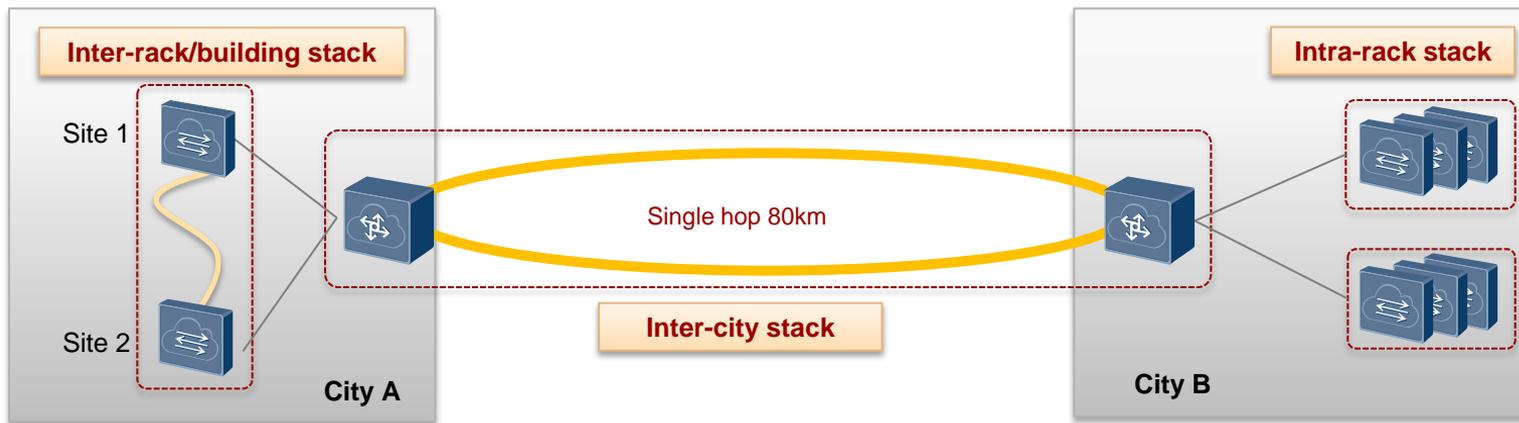
## Improve bandwidth efficiency and reduce CAPEX

- Ring protection protocols such as MSTP are not required, and no link needs to be blocked.
- 100% of bandwidth is used. (Only 50% of bandwidth is used on an STP network).

## Build a highly reliable loop-free network

- CSS, iStack, and Eth-Trunk build a loop-free network.
- The convergence time is much shorter than STP.
- The system is still running when a single device fails.

# Flexible CSS & iStack Deployment



## CSS Stack Capacity

- 2 CE12800 switches in a stack, scalable to 4 (Q2, 2013)
- Stacking with line cards interfaces, up to 640 Gbps stack bandwidth (16\*40 GE), scalable to 1600 Gbps(16\*100 GE)
- Up to 16 CE6800 switches in a stack

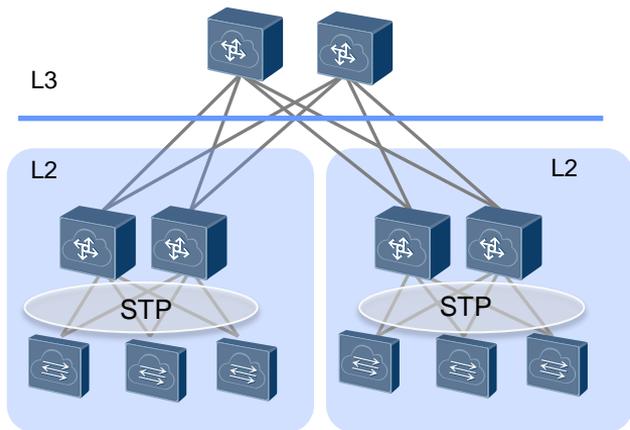
## Long-Distance Stack

- Dedicated stacking cables or common fibers for stacking
- Single-hop distance of 80 km (10 GE SFP+) when using fibers for stacking
- Applicable to various scenarios

## Preferred Local Forwarding

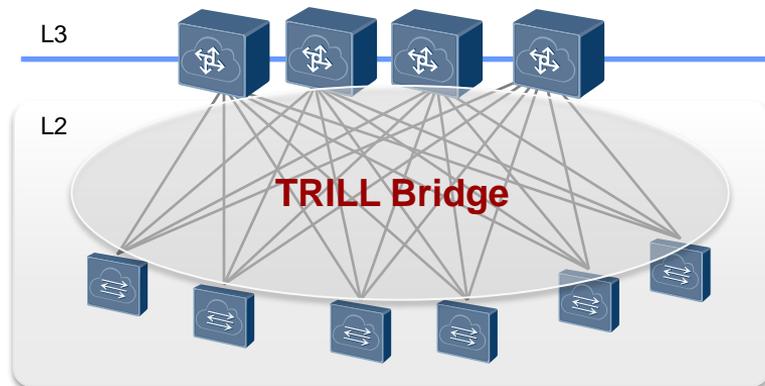
- Load balance traffic among links of cluster members
- Prefers local paths for forwarding
- Improve bandwidth efficiency

# TRILL Builds Next-Generation Data Center



Traditional DC network architecture

- Use STP to build a loop-free network: low bandwidth efficiency
- VMs cannot migrate across a Layer 3 network: deployment restricted

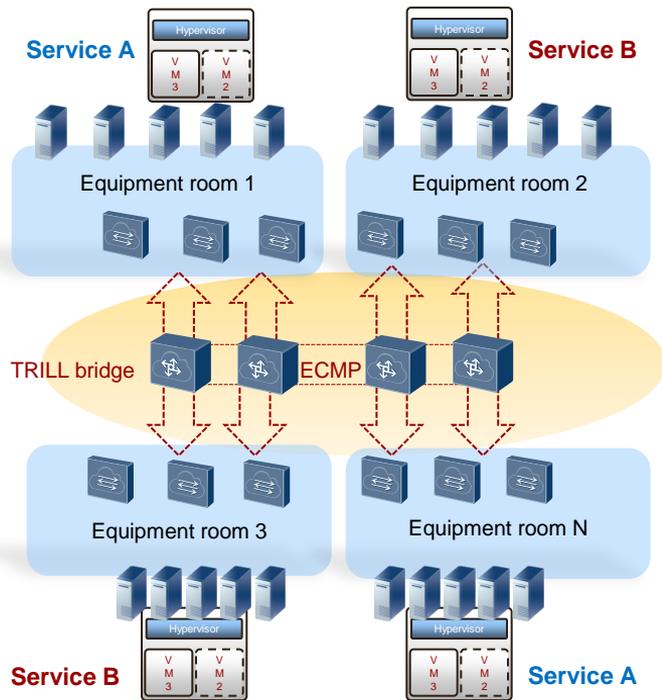


New DC network architecture

- Use TRILL to build a large Layer 2 network: allows VM migration in a wide range
- Supports a maximum of 32 equal-cost links: scalable network

**Standard IETF TRILL protocol, L2 interconnection between multi-vendor devices**

# Benefits of TRILL-based Large L2 Network



## Efficient, Reliable Forwarding

- Data traffic is forwarded along the shortest path through **SPF**.
- **ECMP** realizes fast forwarding over multiple equal-cost links.
- Real-time detection of link/node failures achieves **sub-second convergence** upon topology changes.

## Flexible Service Deployment

- VMs can migrate freely in a large L2 network, making service deployment more flexible.
- A logical network supports 4 K tenants and allows network resources to be allocated on demand.
- Unicast and multicast services are controlled by a single protocol, which is easy to configure.

# Contents

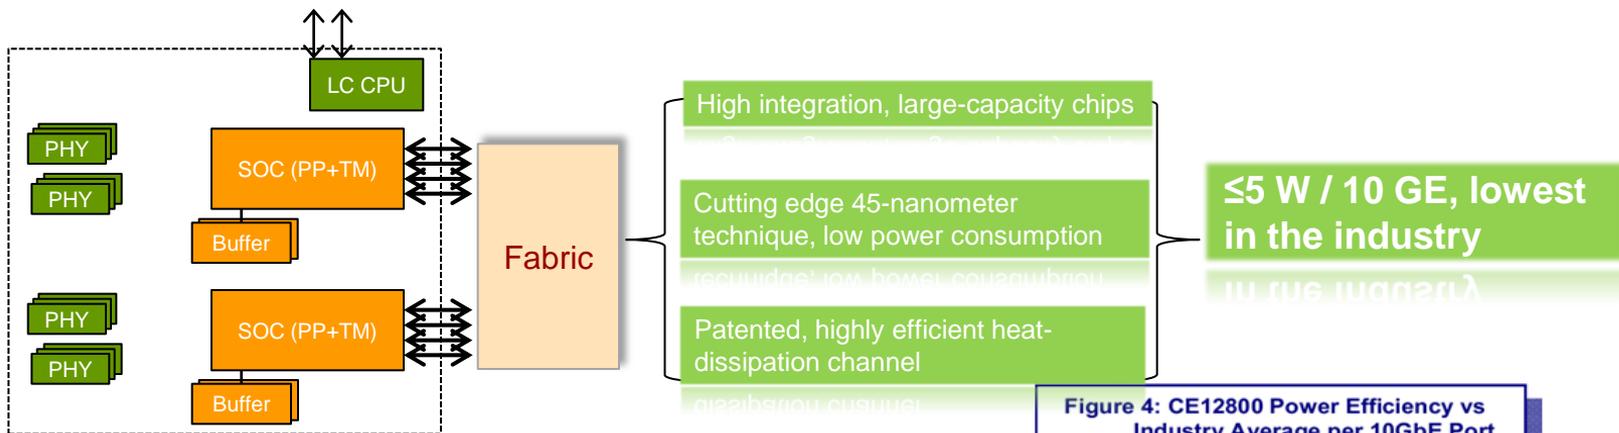
- **Next-Generation Core Switch—CloudEngine 12800**
  - Next-Generation Core Engine to Provide the Highest Performance in the World
  - Advanced Architecture to Enhance Efficiency and Reliability
  - Comprehensive Virtualization Capabilities to Support Flexible Deployment
  - Innovative Energy Saving Technologies to Build a Green Data Center

# Innovative Energy-Saving Technology



- High integration and lowest power consumption
- Patented ventilation channel design, smart heat dissipation
- On-demand power supply, dynamic energy-saving
- Flexible cabling, low noise

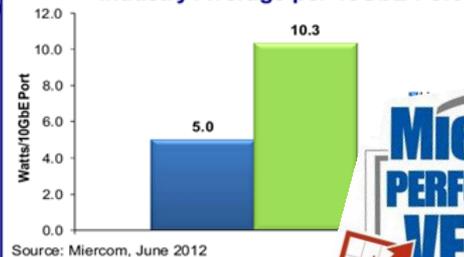
# Lowest Power Consumption per 10GE Port



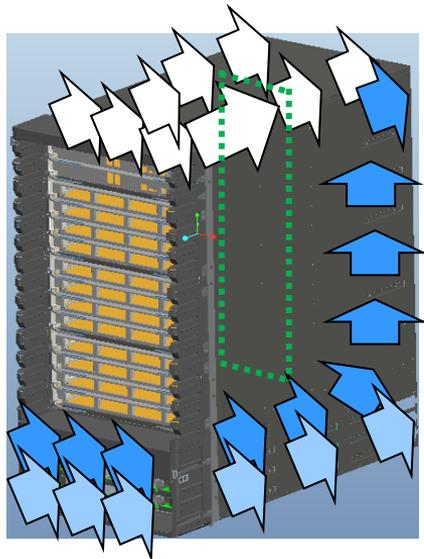
## Key findings and conclusions:

- Highly reliable Core Data Center Switch with hot swappable components including line cards, power supplies, cooling fans and redundant processing modules
- CE12800 is highly energy efficient with power consumption 50% better than the industry average for core switches

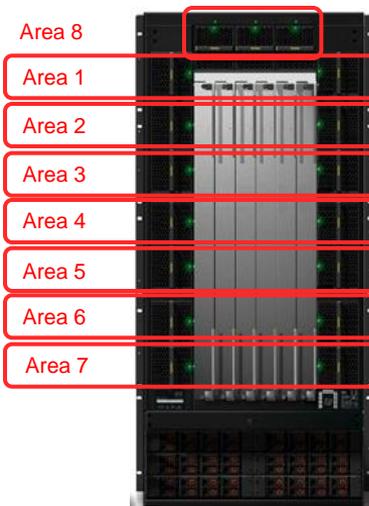
Figure 4: CE12800 Power Efficiency vs Industry Average per 10GbE Port



# Patented Ventilation Channel Design



- Front-to-rear ventilation channel design, which is suitable for parallel cabinet arrangement in a data center equipment room and can reduce power consumption of air conditioners in the equipment room.
- The straight ventilation channels speed up heat dissipation of line cards.

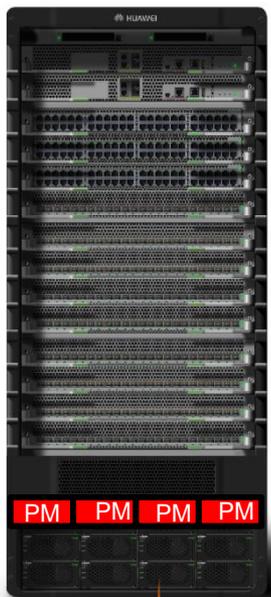


**Double counter-rotating fans**  
**Highly efficient heat dissipation**

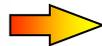
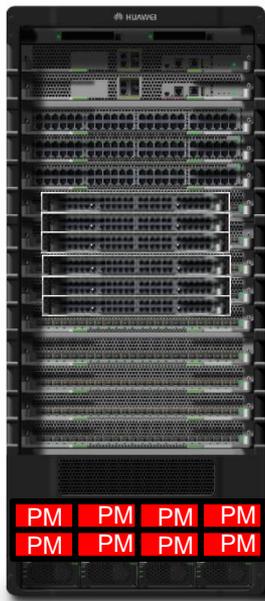
## On-demand smart fan speed adjustment

- A chassis is usually not fully loaded with line cards. If two line cards are configured, only fans in area 1, 2, and 8 work, and fans in other areas do not rotate or rotate at minimum speed to reduce power consumption and noises.
- Since large capacity high-performance line cards (such as 24\*4GE) consume much more power than small capacity line cards (such as 24\*10GE), install them in different areas. The system collects data about temperatures of all areas and adjusts fan speed accordingly. The system improves fan speed in areas where line cards consume more power, while reducing the speed of other fans to cut power consumption and noise.

# Smart Management and On-Demand Deployment



The PMs and power supply can be flexibly increased according to the number and power consumption of line cards.



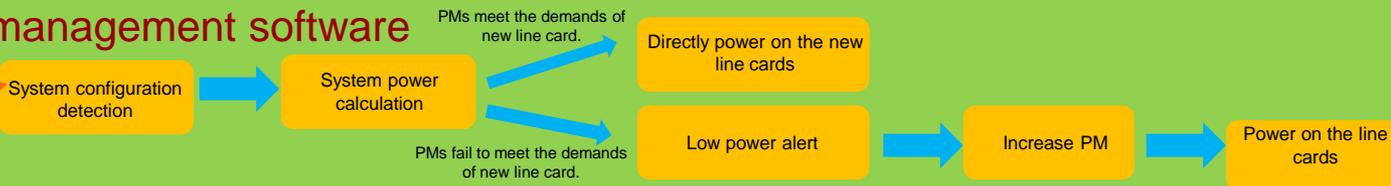
## On-demand deployment of PM

- Deployment of proper number of PMs reduces initial investment.
- Fine-grained control of PMs. The PMs can be increased on demand.

## High power conversion rate, low power conversion loss

- On-demand PM configuration increases load of PMs to over 80 percent and improves power conversion rate from 70 to 96 percent in lightly loaded circumstances, which saves energy.

## Smart power management software



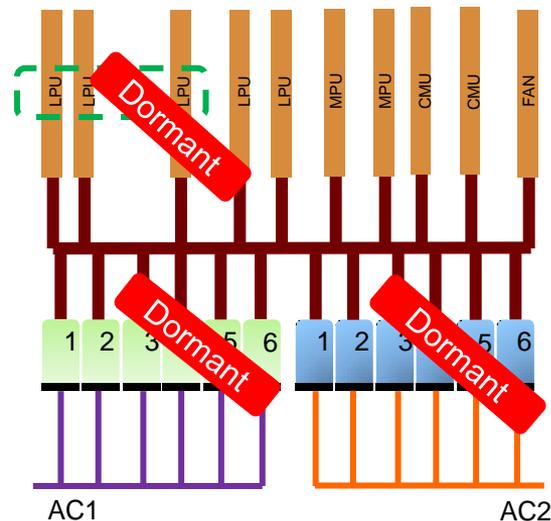
# Realtime Monitoring and Dynamic Energy Saving

## Smart PM control based on service traffic



— Normal power consumption    — Service traffic    - - - Smart control of power consumption based on service traffic

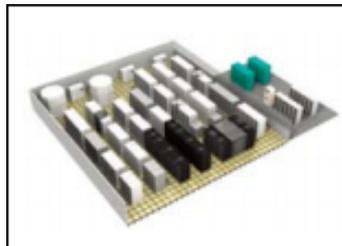
- ① All the devices are running, and the power consumption stays at its peak.
- ② The service traffic fluctuates as time changes.
- ③ The system reduces power consumption and realizes on-demand power supply without affecting service traffic.



- The system measures power consumption in real time and sets one or more power modules in sleep mode when the system workload is low. This improves power efficiency and reduces power consumption.

# Flexible Cabling, Low Noises

1 m deep cabinet facilitates cable layout.

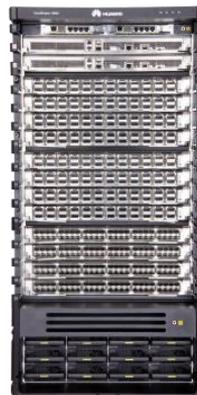


The smallest cabinet is only 1 m deep, which is easy to deploy.



Easy cabling and retractable cable tray

Self-developed noise reduction cabinet reduces noises by 30 dB.



Patented noise reduction technology reduces the switch noise to the level of box switches.

# CloudEngine 6800/5800 Series

— Aggregation switch for data center and high-end campus networks

# CloudEngine 10 GE TOR Switch



CE6850-48T4Q-EI



CE6850-48S4Q-EI

	CE6850-48S4Q-EI	CE6850-48T4Q-EI
Downlink interface	48*10 G SFP+	48*10 G Base-T
Uplink interface	4*40 G QSFP+	4*40 G QSFP+
Switching capability	1.28 Tbps	
Forwarding performance	960 Mpps	
Direction of ventilation channel	Front-to-rear/rear-to-front	
Dimensions (width x depth x height)	442 x 600 x 43.6	
Supports a variety of features for s for data centers, best choice for cloud-computing data centers		

# CloudEngine GE TOR Switch



CE5850-48T4S2Q-EI

40 GE uplink GE TOR switch:

1. Non-blocking 40 GE interconnection network, 1 +1 link backup
2. Stack without shrinking the port bandwidth
  - 1) Stack bandwidth: 120 G (unidirectional). Realizes full mesh stack of three devices without shrinking the port bandwidth.
  - 2) The stack bandwidth can reach 80 G without occupying the 4\*10 GE uplink service interface.

	CE5850-48T4S2Q-EI
Downlink interface	48*10/100/1000 BASE-T
Uplink interface	4*10 G SFP+ + 2*40 G QSFP+
Switching capability	336 Gbps
Forwarding performance	252 Mpps
Direction of ventilation channel	Front-to-rear/rear-to-front channel
Dimensions (width x depth x height)	442 x 420 x 43.6

# High Performance, High Density, 40 GE Uplink



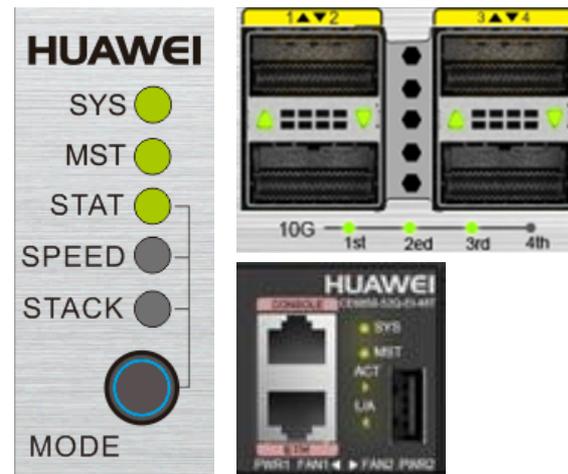
CE6850-48T4Q-EI

- **Among the 1 U high switches**
  - Highest performance: CE5850 - 336 Gbps, CE6850 - 1.28 Tbps
  - High-density access: CE6850 provides 64\*10 GE interfaces. (QSFP+ can be used as 4\*10 GE)
- **CE5850 supports 4\*10 G + 2\*40 G uplink interface:**
  - 40 GE uplink, builds a full 40 G high-performance data center network together with CE12800
  - CE5850: the first access switch that supports 40 GE uplink



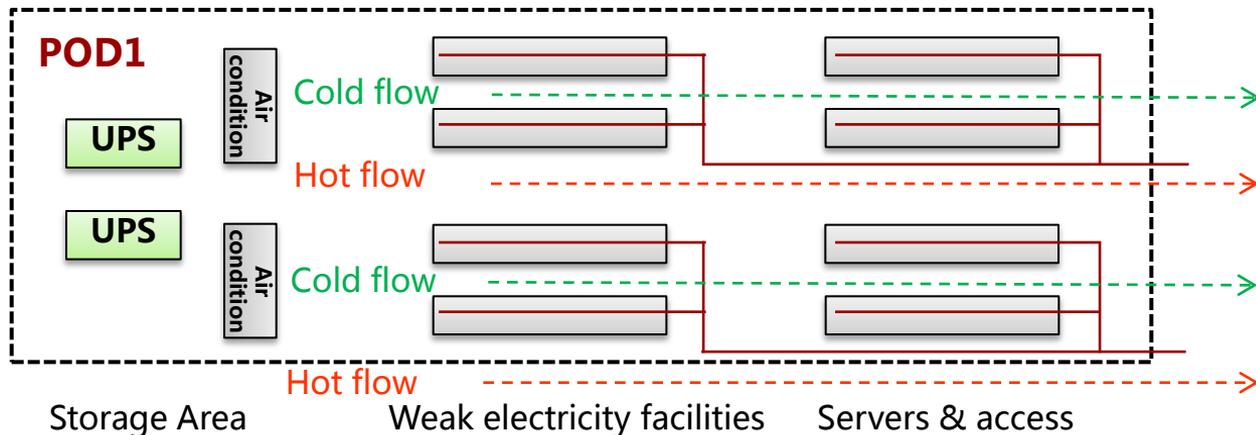
# High Maintainability

- **Easy maintenance:**
  - Management interface, fan, and power module are at the front side of the switch, which are easy to maintain.
  - The data interfaces are at the back of the switch and face servers, making cabling and maintenance easy.
- **Clear indicators:**
  - Indicators on the front and the back panels display the system status and the stack status, which facilitates maintenance.
  - The stack indicator displays information about the role and identity of the device in the stack.
  - Port indicators help to monitor all ports.
    - Creative 40 G port indicator displays status of all sub-ports.
    - Mode switching button enables port indicators to display port status as well as port rate.
  - Remote location
    - The indicator of a device can be remotely turned on by the NMS or the control platform, making it easy to locate the device in the crowded equipment room.
    - The bright blue indicator makes it easy to find the device.



# Front-to-Rear/Rear-to-Front Ventilation Channel

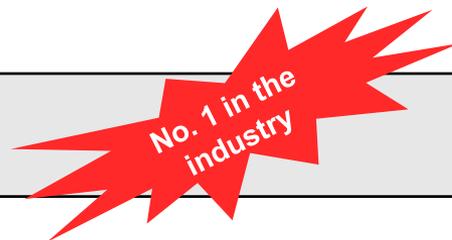
- Flexible selection of fan modules can support front-to-rear/rear-to-front ventilation channel.
  - The fan modules and power modules must be compatible.



The CE5800 and CE6800 are designed with strict front-to-rear ventilation channels that can separate hot air flows from cold air flows without transforming device structure. This design improves heat dissipation efficiency and reduces power consumption.

# High-Specification TOR

	<b>CE6850</b>	<b>CE5850</b>
MAC table capacity	128 K	128 K (large MAC mode)
FIB capacity	16 K	16 K
Multicast replication capability	32 K	32 K
Number of TRILL nodes	512	512



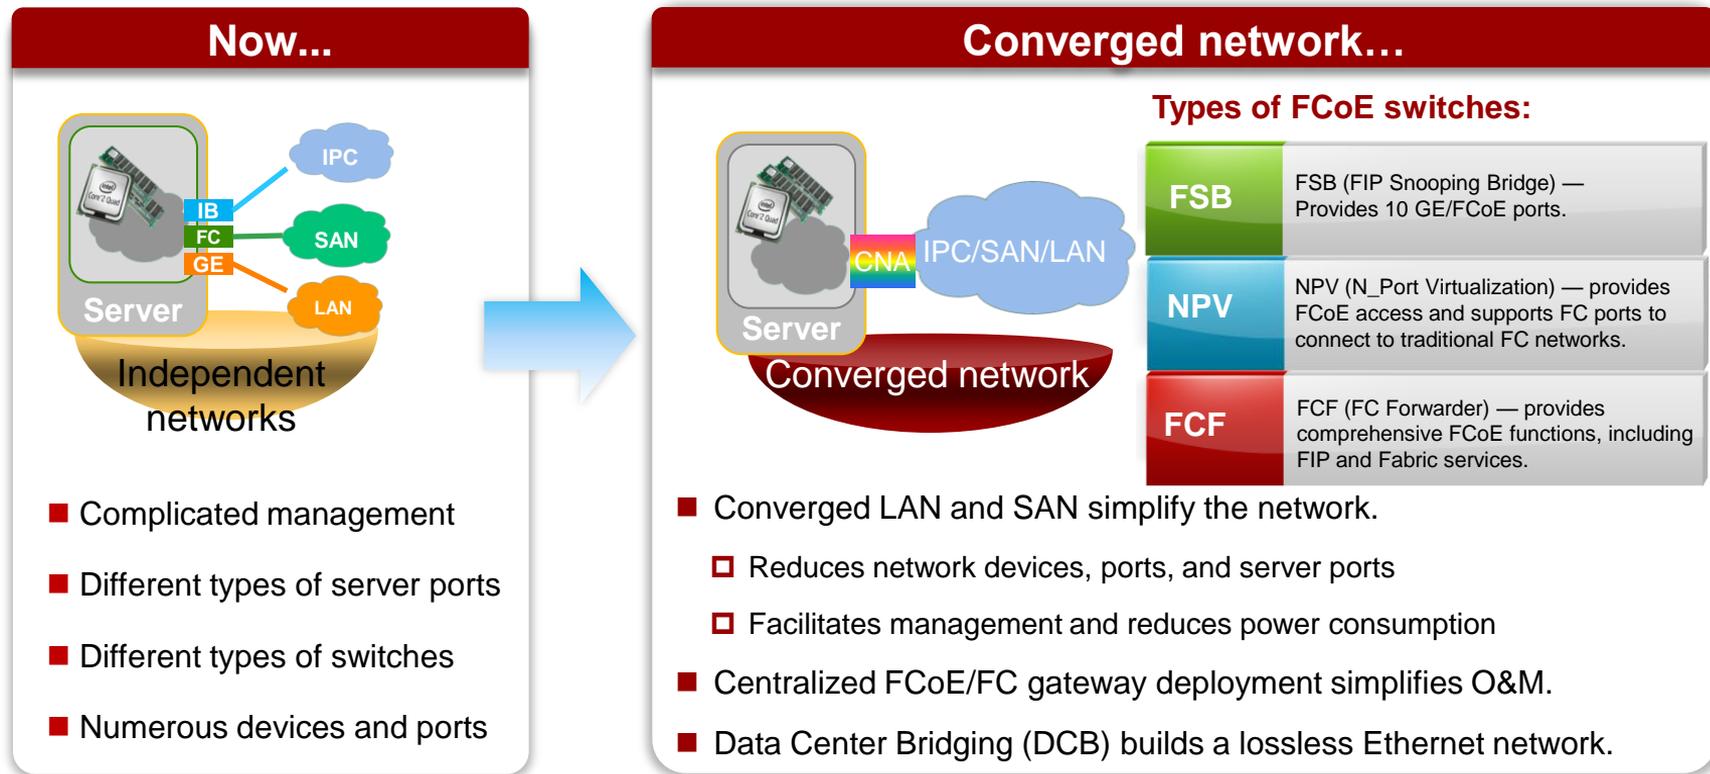
- ❑ Realizes effective virtual aggregation.
- ❑ Supports large-capacity stack system.

# High-Performance Stack

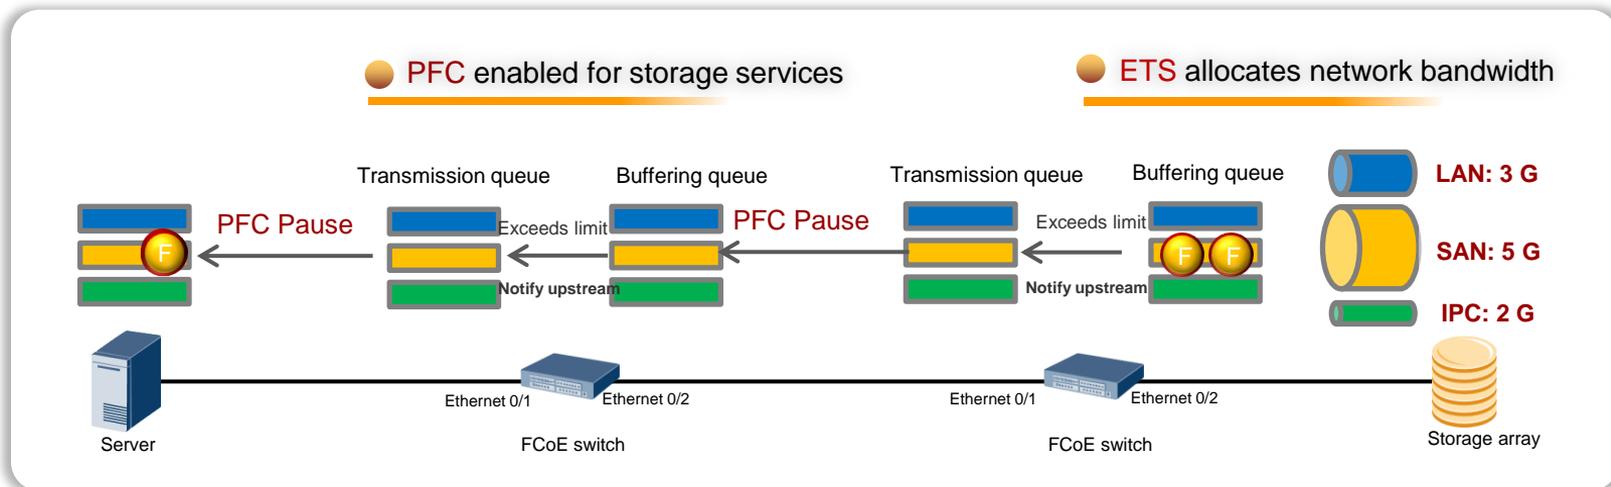
- **High-performance stack system**
  - A maximum of 16 devices in a stack.
  - More than 768 GE/10 GE access interfaces in a stack.
- **High stack bandwidth**
  - The CE5800 supports 120 G stack bandwidth, and the CE6800 supports 160 G.
- **Quick recovery**
  - Recovery in 200 ms
- **Easy to operate**
  - Supports pre-design deployment and offline configuration. New devices and components can be added and immediately used, realizing "pay as you grow." Supports ID setting of devices in the stack system, facilitating device location and maintenance.
  - Indicators clearly display the role and status of the stack system, realizing basic stack system maintenance without a computer.
  - Supports quick and automatic software upgrades of the stack system, simplifying the upgrade process and reducing the required work.



# FCoE Builds a Converged Network



# DCB Guarantees Zero Packet Loss

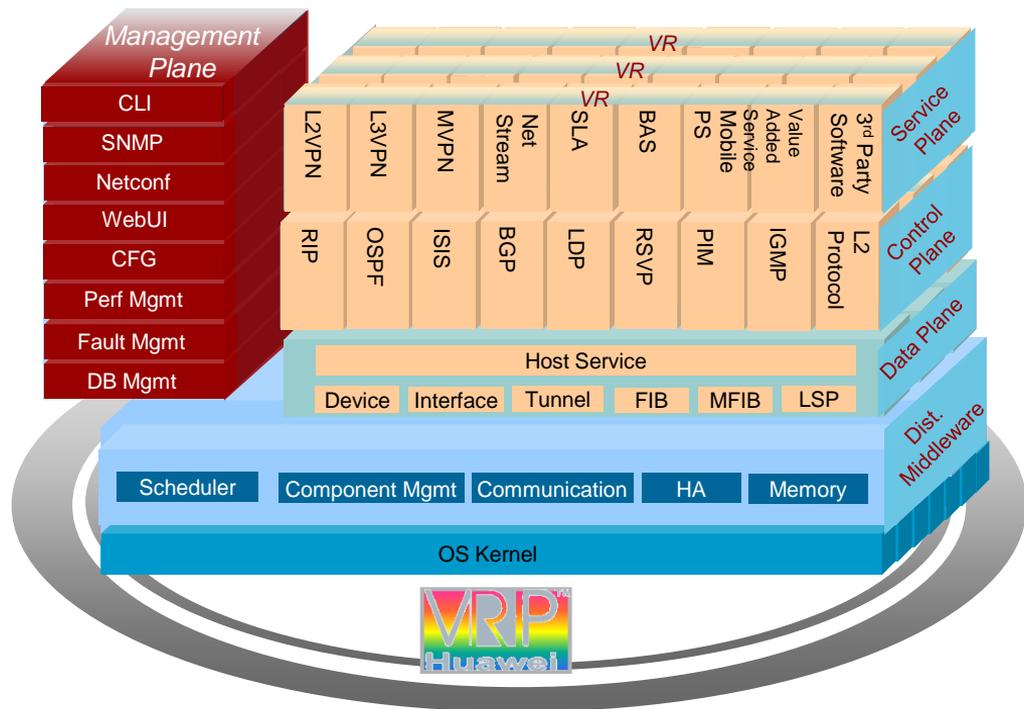


- **DCBX (Data Center Bridging Exchange):**  
Discovers peer device and negotiates PFC parameters
- **PFC(Priority-Based Flow Control) 802.1Qbb :**  
Sends a Pause frame request the upstream device to stop sending data when the traffic rate exceeds the limit.
- **ETS (Enhanced Transmission Selection) 802.1Qaz :**  
Allocates bandwidth for each type of data, for example, allocate 3 Gbps, 5 Gbps, and 2 Gbps bandwidth to LAN, SAN, and IPC services respectively on a 10 GE link.

# VRP8 Software System Platform

— Exclusively for CloudEngine series switches

# Stable Software Platform VRP8



## Stable

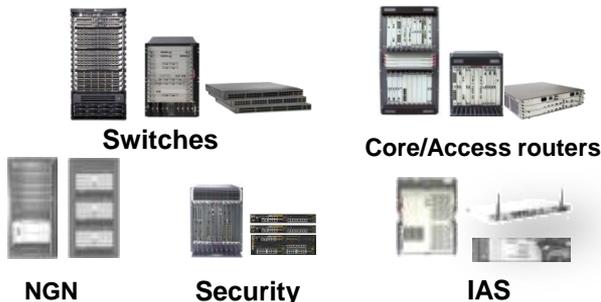
- Applies to **five** Huawei product families, serves users in over **100** countries and regions, and has been verified by stable operation of a huge number of products.
- Non-Stop Service for Anything (NSA) ensures high service availability.
- Provides a versatile platform for core and TOR switches, simplifying maintenance and reducing maintenance workload.

## High Performance

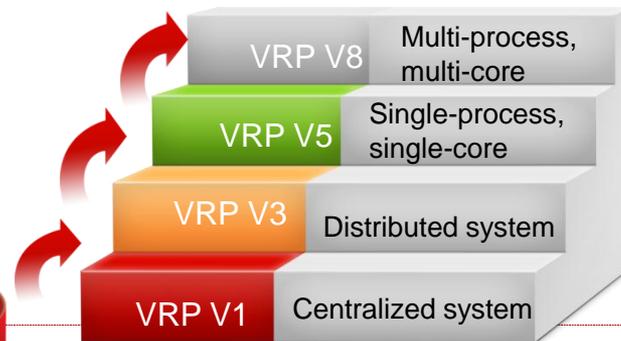
- **Distributed multi-service system** deploys a single protocol or service on multiple CPUs or CPU cores for distributed processing, **improving service processing performance**.
- System performance is not limited by capability of a single core or memory of a single CPU, and faults can be located more precisely. **Service reliability is improved**.
- Multi-core/multi-CPU processing maximizes system performance **and allocates resources on demand in multi-tenant applications**.

# Verified Stability

Applied to **5** product families



**18** years of experience



**VRP**

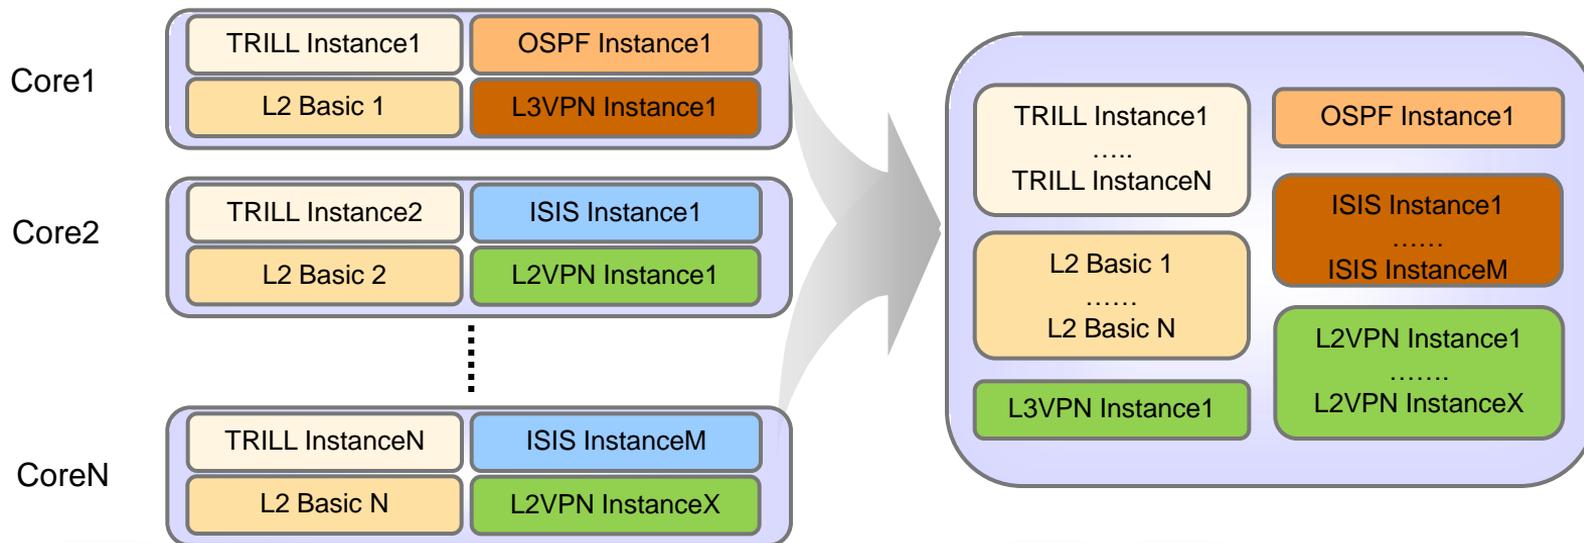


Serves **100+** countries and regions,  
stable operations of **7,000,000+**  
devices

**2000+** R&D engineers dedicated in  
innovations  
Active participant in standardization  
organizations



# High Performance and Reliability



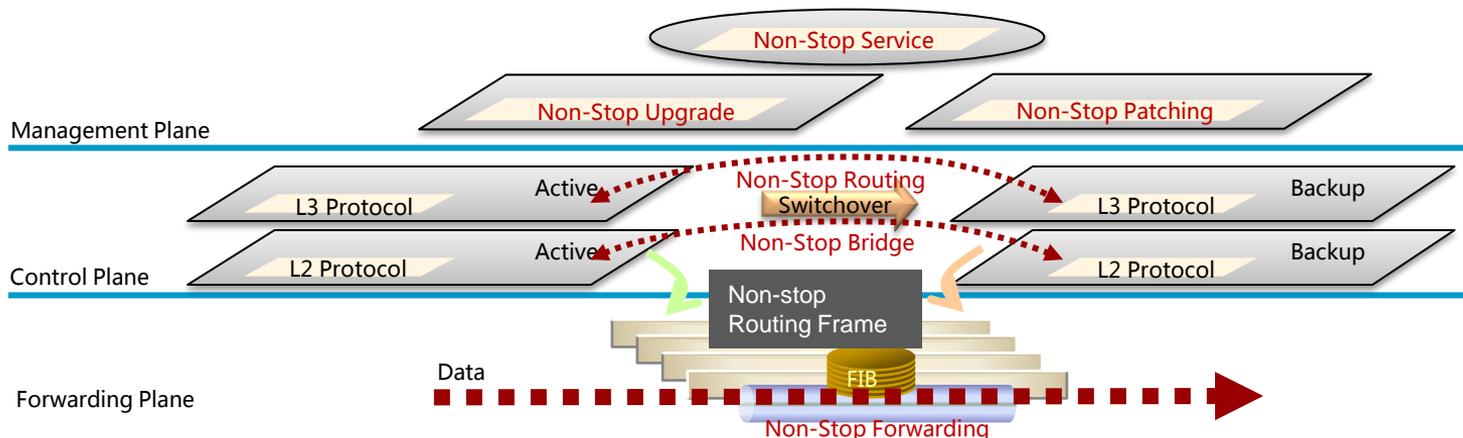
## High performance

- Parallel computing on multiple cores for one service, significantly improving service processing performance
- Different services on different cores, meeting diversified requirements for service processing

## High reliability

- Independent processes—one process does not affect another
- Modular design—isolates memory space and improves system reliability

# Whole Lifecycle Reliability System——NSA



Function	Scenario	Description
NSF	Forwarding plane	Non-stop service forwarding capability
NSR/NSB	L2/L3 control plane	Non-stop capability in layer 2/3 control plane
NSU	System upgrade	Non-stop service capability while system upgrades
NSP	System Patching	Non-stop service capability while system upgrades
NSS	Service	Keep service online

**VRPV8 provides whole lifecycle non-stop recovery technologies NSA (Non-Stop Anything), to guarantee non-stop service in switchover.**

# Thank you

[www.huawei.com](http://www.huawei.com)

**Copyright©2011 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.