



Huawei Launches CloudFabric Solution to Build an Ultra-simplified and Efficient Cloud Data Center Network

Overview

Solution background: Digital transformation of industries accelerates the use of cloud-based data centers by enterprises. It is predicted that by 2025, 85 percent of enterprise applications will be deployed on the cloud. As the basis of building the cloud, the data center network is the key factor in digital transformation, but faces the following challenges:

- Complex architecture design: Thousands of products from over 10 vendors are available for customers, and the solution design period often lasts for months.
- Long deployment period: Several weeks are required due to long process involving requirement analysis, manual configuration and commissioning.
- Difficult fault locating: In the past ten years, managed objects have increased exponentially, which brings great pressure on network operations. Without visibility capabilities, networks are just like black boxes, and network faults are hard to locate.

Solution overview: Huawei has launched the CloudFabric solution to redefine network design, deployment, operations, and interconnection of data centers, aiming to build a simplified cloud data center network. The CloudFabric solution are optimized for specific environments, starting with High-Availability/Hyperscale/HPC. CloudFabric for High-Availability, CloudFabric for Hyperscale, and CloudFabric for high performance computing (HPC) cover design, deployment, and operations, realizing fast and flexible solution design, automatic network deployment, and intelligent network operations.

- Simplified design: CloudFabric support on-demand and self-service customization in different scenarios, helping to quickly design a cloud data center network solution.
- Simplified deployment: Agile Controller provides drag-and-drop and automatic service orchestration, shortening the service provisioning time from weeks to minutes.
- Simplified operations: Fabric Insight, an intelligent network analysis platform enables network visibility from the perspective of applications and helps locate faults within minutes.

Media Roundtable Agenda

Topic: Huawei Launches CloudFabric Solution to Build an Ultra-simplified and Efficient Cloud Data Center Network

Time: September 6, 2017 15:30-16:15

Location: Function Room 2, 3rd floor, Kerry Hotel

Speakers: Ian Foo (Director of Network Product Line Product Marketing & Branding, Huawei Technologies Co., Ltd.);

Yu Tao (Director of Data Center Network Solution, Huawei Technologies Co., Ltd.);

Ma Ye (Vice President of the Router & Carrier Ethernet Product Line, Huawei Technologies Co., Ltd.);

Schedule (45 minutes):

Time	Topic	Speaker
------	-------	---------



15:30 – 15:45	Huawei Launches the CloudFabric Solution Optimized for Specific Environments, Starting with High Availability/Hyperscale/HPC	Ian Foo
15:45 – 16:15	Media Q&A	Three speakers: Ian Foo, Yu Tao, and Ma Ye

Key Messages (Solution’s Differentiation Capabilities)

CloudFabric for High-Availability

1. Current Status and Challenges of Availability-Sensitive Data Center Networks

For some verticals requiring high availability, such as financial, data center networks are faced with the following challenges:

- **Reliability:** Global banking regulators require that financial services be provided even if the financial IT system fails or a devastating disaster occurs. For example, the China Banking Regulatory Commission requires that if a bank network is interrupted for over 3 hours, a Level-1 fault should be reported to the State Council of the People’s Republic of China.
- **Security:** The data center uses zero-trust network security to ensure east-west traffic security. Existing networks of banks experience thousands of complex, manual security policy adjustments each year, which are error-prone and can be highly impactful.
- **Operations:** Service development brings great pressure to network operations. On the one hand, the network management scale becomes larger with the development of cloud-based data centers. Using a financial enterprise as an example, in the past five years, the number of network elements increased 10 times and the number of faults increased 18 times. On the other hand, network faults must be located and recovered in the shortest time because bank services on live networks have pressing priorities.

2. Capabilities of CloudFabric for High-Availability

- **Highly reliable network and continuous service:** (1) Link-level: Hardware-based Bi-directional Forwarding Detection (BFD) technology detects link status in milliseconds. (2) Device-level: The CE12800 core switches provide 99.999% carrier-class reliability and TOR switches provide a Multi-chassis Link Aggregation Group (M-LAG) backup mechanism, ensuring continuous services.
- **High security for applications and precise management and control for security policies:** CloudFabric supports micro-segmentation. Security policies, which are configured by subnet originally, can be configured by VM to ensure east-west traffic security between applications. Security policies can be automatically deployed within minutes by GUI-based drag-and-drop operations.
- **Intelligent operations and fault location within minutes:** Based on the collected data flows from the entire network, Fabric Insight analyzes and displays the application and network health status in real time and actively predicts risks. A Big Data platform supports data searched among tens of billions of historical records and analysis of millions of packets within seconds, realizing precise playback of historical faults. In the event of application exceptions, applications are automatically associated with faulty links, implementing fault location within minutes.

3. Partners

- Cloud platform layer: Agile Controller can interoperate with cloud platforms FusionSphere, VMware, Microsoft, OpenStack, Red Hat, UnitedStack, CloudStack, EasyStack, achieving unified management of networks and IT resources.



- Controller and management tool layer: Agile controller is based on the ONOS open source platform. CloudFabric supports VMware NSX and Brocade Vyatta, achieving unified scheduling and automatic operations management of services. Huawei has jointly published the data center network automatic operations solution with Puppet Enterprise (U.S.), making development 30 times more efficient. Huawei's intelligent network analysis platform, Fabric Insight, cooperates with a system of Tingyun, an industry-leading Application Performance Management (APM) solution provider, to support associations between applications and networks. Application exceptions can be mapped to network links with one click, implementing E2E operations and network evaluation.
- Network device layer:
 - Agile Controller, based on open APIs and southbound interfaces such as OpenFlow, seamlessly integrates with applications from F5 Networks, an industry-leading applications delivery vendor, to jointly release the active-active data center solution.
 - The Agile Controller provides the DNS, DHCP, and IPAM integrated management solution together with Infoblox, achieving automatic management of application policies.
 - The Agile controller unites with CheckPoint to create a secure and reliable data center security solution.
 - The Agile Controller cooperates with Intel, Mellanox, Broadcom, and Qlogic in the hardware and chip fields, creating a highly capable data center network.
- Computing virtualization layer: CloudFabric support mainstream computing virtualization platforms such as VMware, Microsoft, and KVM to achieve unified linkage between a data center's physical and virtual networks.

CloudFabric for Hyperscale

1. Challenges of Hyperscale Data Center Networks

For some verticals, such as ISP, data center networks continue to expand, posing the following challenges:

- **Pressure of a scalable data center and network performance:** ISPs' data centers are expanding and some super-large data centers are deployed with ten thousands of servers. In the spine-leaf architecture, core switches and TOR switches are fully meshed, causing a large number of cable connections and complex management. Video-related and search-related Big Data service traffic is increasing rapidly, and multiple servers often attempt to access one other server simultaneously, resulting in burst traffic that exceeds the available bandwidth.
- **Pressure of automatic network deployment:** In a data center with over ten thousand servers and thousands of network devices, a large number of duplicate network device configurations need to be automatically completed by scripts. At the same time, the network changes from a platform for supporting services to a commercial off-the-shelf cloud offering that can be sold to an enterprise. The network must be deployed within minutes, requiring greater network automation.
- **Pressure of multi-cloud collaboration and management:** Different Internet services have different requirements for I/O performance of computing resources. Data center networks need to support access and unified cloud-based management of multiple types of computing resources, including bare metal servers, Virtual Machines (VMs), and containers. To unify management of computing and storage networks, data center networks need to interconnect with mainstream cloud management platforms from different vendors to provide flexible multi-cloud management capability.

2. Capabilities of CloudFabric for Hyperscale

- **Non-blocking access of 50,000 10 GE servers and E2E huge buffer with 80 times the industry average:** The CE12800 core switches provide a 1,032 Tbit/s switching capacity (highest in the industry), and non-blocking access of up to 50,000 servers using a Clos switching architecture and cell-switching technology. The industry's first 100 GE short-distance single-pair 100GE optical transceiver reduces cable connections by 75 percent and simplifies management. Line cards of CE12800 core switches and CE6870 ToR switches support 24 GB and 4 GB buffers, respectively, making them well suited to cope with burst traffic on the network.

- **Fully automatic network and service provisioning within minutes:** Huawei' Agile Controller provides drag-and-drop and automatic service orchestration, shortening the service provisioning time from weeks to minutes. The CloudEngine switches can be managed by Devops tools (such as Ansible and Puppet) to automatically deliver configurations to the underlay network.
- **Multi-cloud collaboration, reducing OPEX by 80 percent:** Data center networks support access and unified cloud-based management of computing resources such as bare metal, VMs, and containers. Huawei's SDN-based Agile Controller can interoperate with cloud platforms of FusionSphere, VMware, Microsoft, OpenStack, and other mainstream platform manufacturers and standard organizations, achieving unified management of network and IT resources. The Agile Controller also supports unified management of multiple data centers, public clouds, and private clouds.

3. Partners

- Cloud platform layer: Agile Controller can interoperate with cloud platforms FusionSphere, VMware, Microsoft, OpenStack, Red Hat, UnitedStack, CloudStack, EasyStack, achieving unified management of networks and IT resources.
- Controller and management tool layer: Agile controller is based on the ONOS open source platform. CloudFabric supports VMware NSX and Brocade Vyatta, achieving unified scheduling and automatic operations management of services. Huawei has jointly published the data center network automatic operations solution with Puppet Enterprise (U.S.), making development 30 times more efficient. Huawei's intelligent network analysis platform, Fabric Insight, cooperates with a system of Tingyun, an industry-leading Application Performance Management (APM) solution provider, to support associations between applications and networks. Application exceptions can be mapped to network links with one click, implementing E2E operations and network evaluation.
- Network device layer:
 - Agile Controller, based on open APIs and southbound interfaces such as OpenFlow, seamlessly integrates with applications from F5 Networks, an industry-leading applications delivery vendor, to jointly release the active-active data center solution.
 - The Agile Controller provides the DNS, DHCP, and IPAM integrated management solution together with Infoblox, achieving automatic management of application policies.
 - The Agile controller unites with CheckPoint to create a secure and reliable data center security solution.
 - The Agile Controller cooperates with Intel, Mellanox, Broadcom, and Qlogic in the hardware and chip fields, creating a highly capable data center network.
- Computing virtualization layer: CloudFabric support mainstream computing virtualization platforms such as VMware, Microsoft, and KVM to achieve unified linkage between a data center's physical and virtual networks.

CloudFabric for HPC

1. What Is HPC?

High-Performance Computing (HPC) is a type of special computing, in which a large number of parallel supercomputer clusters closely communicate with each other over networks. An HPC system is able to calculate quickly and accurately, which is helpful to solve problems in the business and science sectors. Therefore, more and more scientific institutions, governments, and enterprises use HPC systems.

2. Challenges of HPC Networks

- **Zero packet loss and low latency required by a computing network:** As an HPC network needs to support dense cluster nodes and high-speed computing, the network requires zero packet loss and low latency. InfiniBand high-speed architecture is widely used currently, which ensures lower latency than traditional Ethernet architecture.
- **Unable to build a large-sized network:** As HPC combines with the business (especially the emergence of AI), HPC networks are expanded from hundreds to thousands of servers, creating a bottleneck in

InfiniBand architecture. Due to complex congestion control algorithms such as credit back-pressure, the InfiniBand architecture supports less than 1,000 servers on a given network, which cannot be expanded elastically.

- **High operations costs due to multiple network structures:** On a traditional HPC network, data transmission and network management use the Ethernet architecture, while the computing network uses the InfiniBand architecture, resulting in high costs for operations from different personnel.

3. Capabilities of CloudFabric for HPC

By leveraging unique innovative algorithms, Huawei launches an Ethernet-based HPC network solution, which offers an equivalent performance with InfiniBand:

- **Low latency and zero packet loss:** Huawei's unique congestion scheduling algorithms of Virtual Input Queue (VIQ) and dynamic Explicit Congestion Notification (ECN) enable traffic scheduling, load balancing, and zero packet loss. The transmission latency is reduced by 90 percent and HPC service computing time is reduced by 40 percent. The overall performance of Huawei's HPC network solution is equal to that of InfiniBand architecture.
- **High networking capability of 5,000 servers:** A spine-leaf 100 GE Ethernet network provides access to 5,000 servers, which is five times that of a dedicated InfiniBand network.
- **Unified architecture management:** Huawei's solution unifies the architecture of transmission, management, and computing networks. Huawei's SDN-based Agile Controller helps build a unified network resource pool that can be scheduled flexibly, making HPC a service.

4. Partners

- Cloud platform layer: Agile Controller can interoperate with cloud platforms FusionSphere, VMware, Microsoft, OpenStack, Red Hat, UnitedStack, CloudStack, EasyStack, achieving unified management of networks and IT resources.
- Controller and management tool layer: Agile controller is based on the ONOS open source platform. CloudFabric supports VMware NSX and Brocade Vyatta, achieving unified scheduling and automatic operations management of services. Huawei has jointly published the data center network automatic operations solution with Puppet Enterprise (U.S.), making development 30 times more efficient. Huawei's intelligent network analysis platform, Fabric Insight, cooperates with a system of Tingyun, an industry-leading Application Performance Management (APM) solution provider, to support associations between applications and networks. Application exceptions can be mapped to network links with one click, implementing E2E operations and network evaluation.
- Network device layer:
 - Agile Controller, based on open APIs and southbound interfaces such as OpenFlow, seamlessly integrates with applications from F5 Networks, an industry-leading applications delivery vendor, to jointly release the active-active data center solution.
 - The Agile Controller provides the DNS, DHCP, and IPAM integrated management solution together with Infoblox, achieving automatic management of application policies.
 - The Agile controller unites with CheckPoint to create a secure and reliable data center security solution.
 - The Agile Controller cooperates with Intel, Mellanox, Broadcom, and Qlogic in the hardware and chip fields, creating a highly capable data center network.
- Computing virtualization layer: CloudFabric support mainstream computing virtualization platforms such as VMware, Microsoft, and KVM to achieve unified linkage between a data center's physical and virtual networks.

Solution: The Best Choice for Interconnecting Cloud Data Centers

1 Current Status and Challenges of Cloud Data Center Networks

As cloud-based Internet services develop, data centers working in active-active or disaster recovery mode are required, posing the following challenges:

- **Cloud-based Internet services and digital transformation result in rapid growth of traffic.** It is predicted that by 2025, 85 percent of enterprise applications will be deployed on the cloud. Inter-DC traffic will exceed 1 Tbit/s. Traffic increases at an average annual growth rate of 30 percent, and data traffic of OTT providers increases at 150 percent. Networks must have large capacities to meet the requirements of rapid service growth, elastic network expansion, and controllable traffic.
- **Slow deployment of traditional network services results in low network utilization.** Network devices from various vendors at different layers cannot be managed and controlled uniformly. Low communication efficiency of different departments greatly affects development of new services. For emergent hot events and a large number of temporary services, traditional distributed route-based forwarding and shortest path policies can easily result in traffic congestion on the network, causing low network utilization.
- **Traditional network architectures and devices that do not suit future data center environments impose high operations costs:** The traditional IP layer is independent of the optical layer, requiring operations by different personnel. High-power-consuming traditional network devices do not suit the data center environment, resulting in high operations costs.

2 Huawei 400G Data Center Interconnection Solution

Huawei's 400G data center interconnection solution uses the industry's first 4 Tbit/s line cards, meeting the rapid growth requirements of inter-DC traffic, simplifying interconnection networks, and reducing OPEX by 75 percent.

- **High bandwidth:** In the solution, a large-capacity router offers a 4 Tbit/s capacity per slot by using a 40-port 100 GE or 8-port 400 GE line card. The port density of such line cards is four times that of a card with 100 GE ports, simplifying optical fiber connections on the backbone network and reducing operations costs.
- **On-demand traffic adjustment:** Huawei's traffic optimization solution uses the innovative ROAM algorithm to centrally calculate path bandwidth and latency; this algorithm is 10 percent faster than others used in the industry. This solution realizes visualized E2E traffic and service detection and adjustment in real time, increasing bandwidth utilization from 30 percent to 80 percent.
- **Energy saving:** Huawei routers use the next-generation Huawei-developed solar 5.0 chips that integrate with NP, MAC, and search algorithms. The high level of integration made possible by 16-nm technology allows the chip to consume 32 percent less power than the previous-generation solution. With the industry's first special card heat dissipation technology, the overall power consumption reduces 50 percent compared with the industry level.