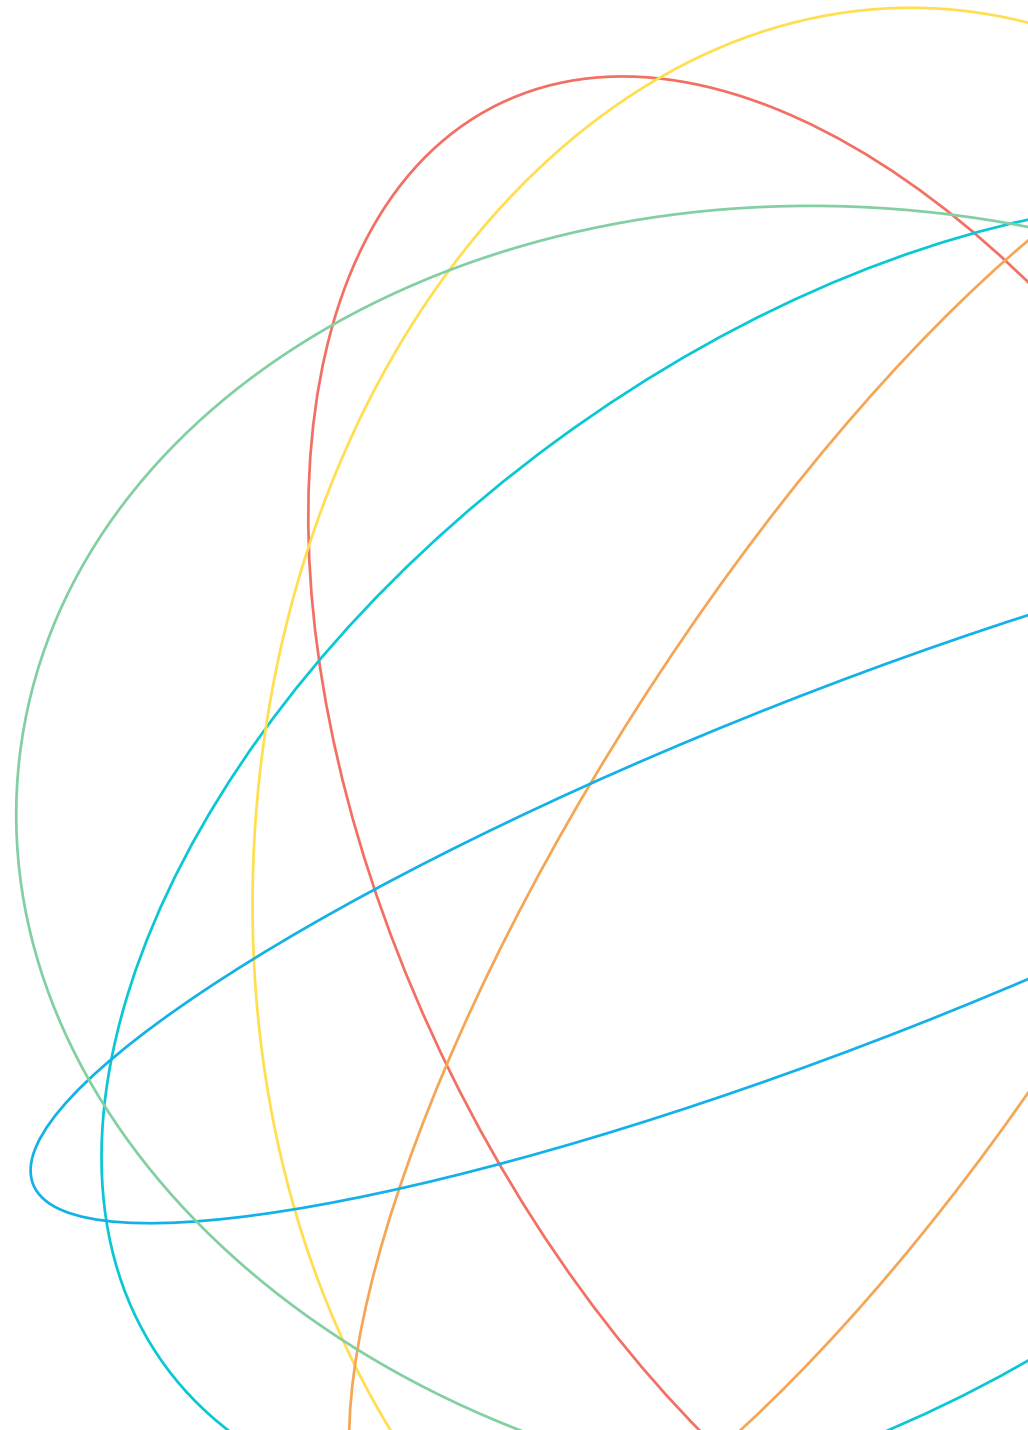
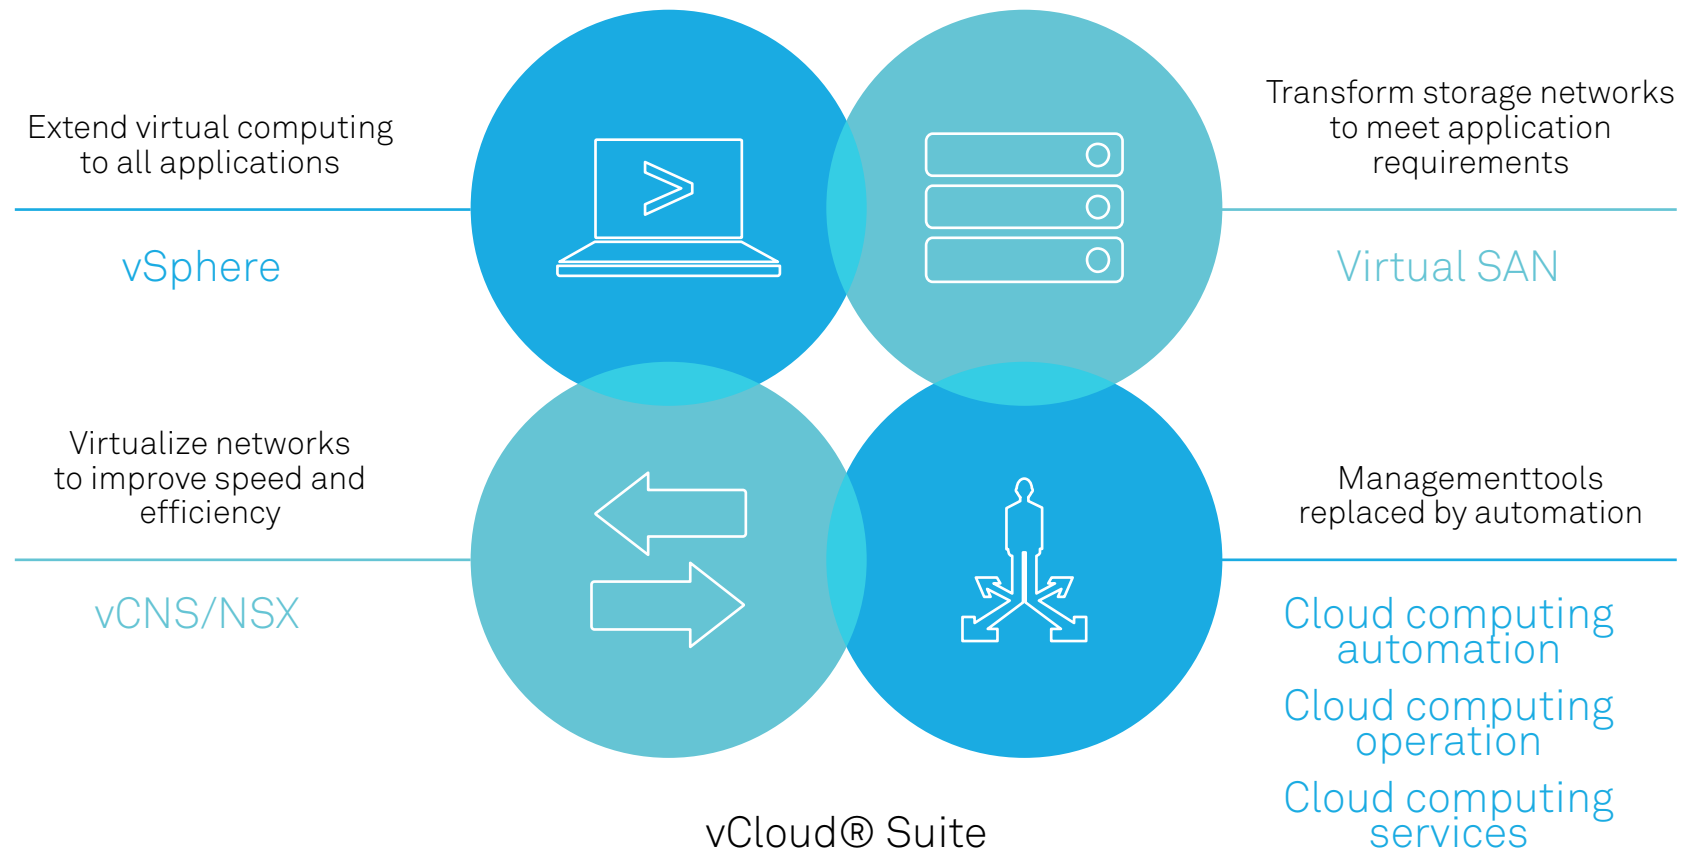




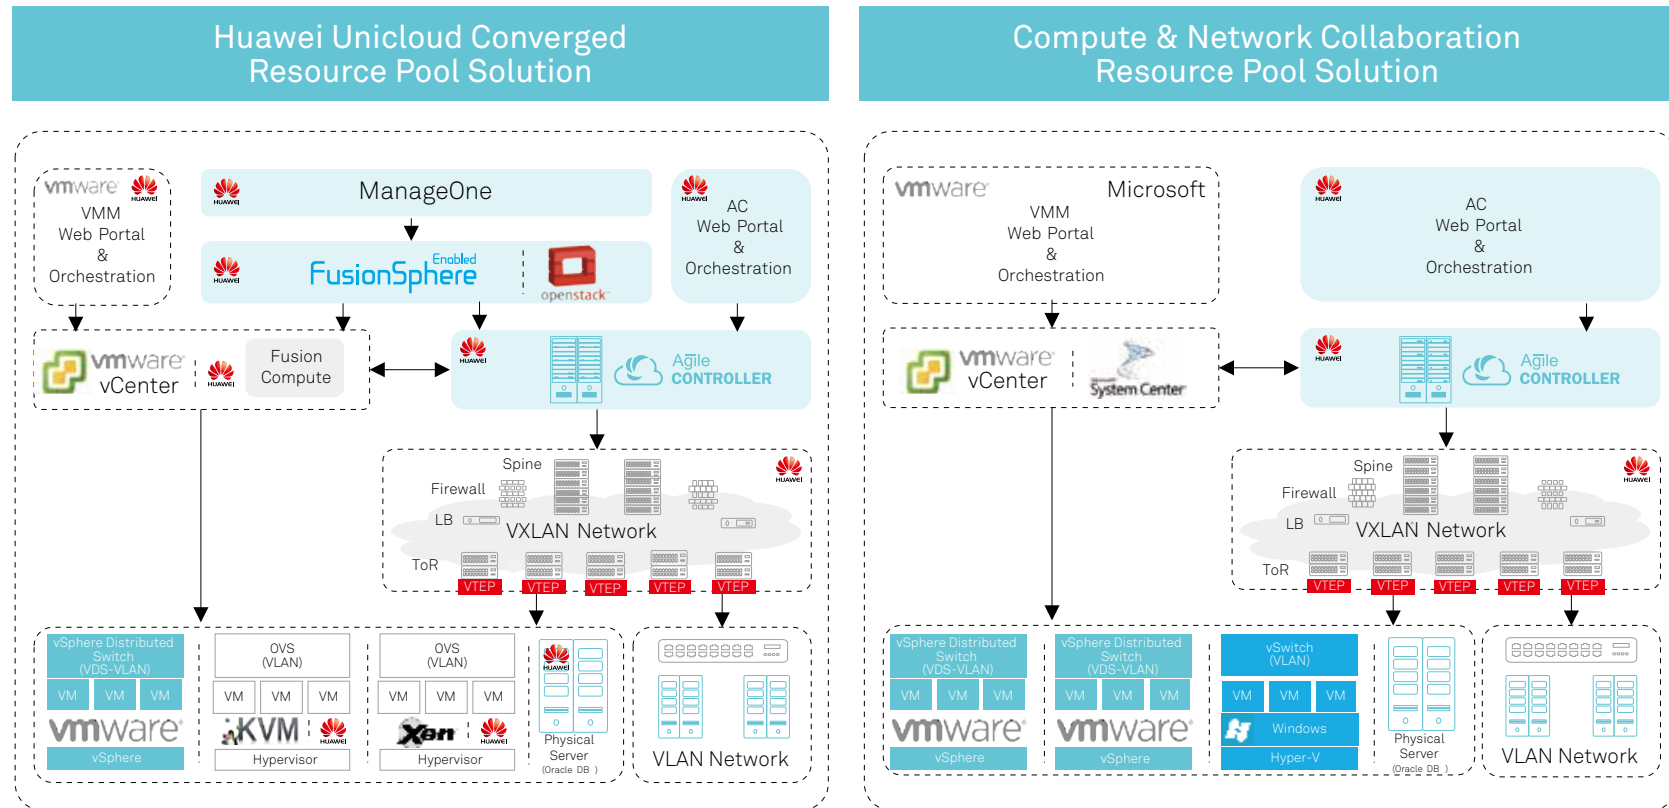
Huawei Cloud Fabric and VMware Collaboration— Innovation Solution in Data Centers



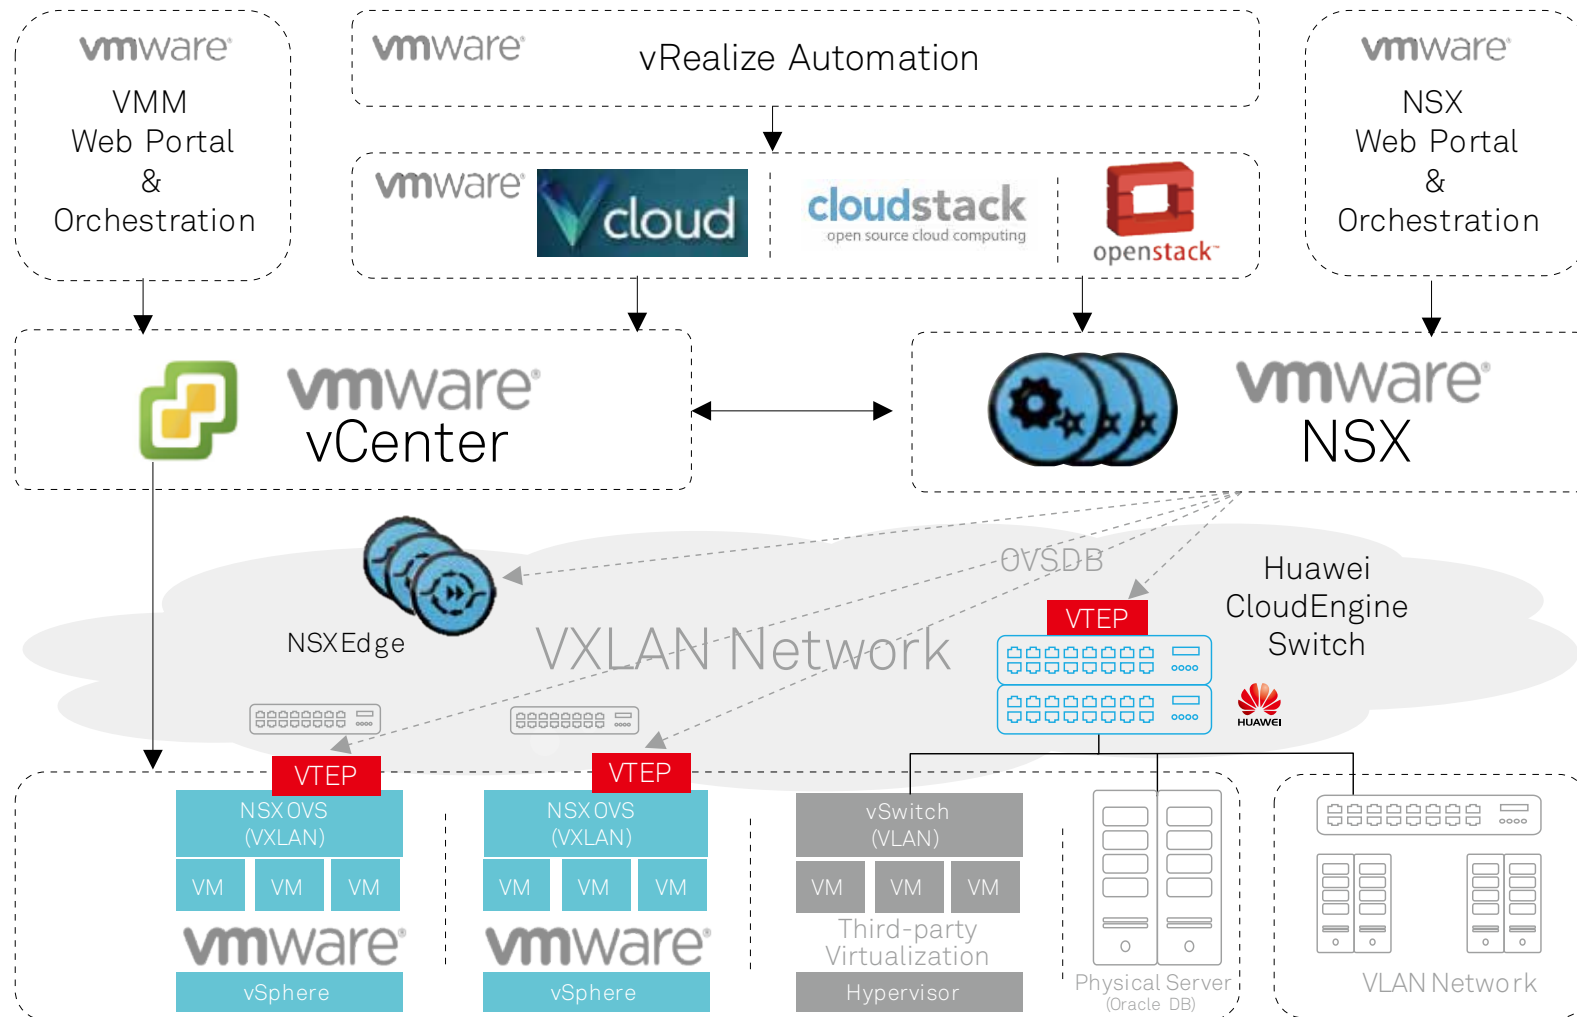
VMware Data Center and Cloud Computing Solution Components



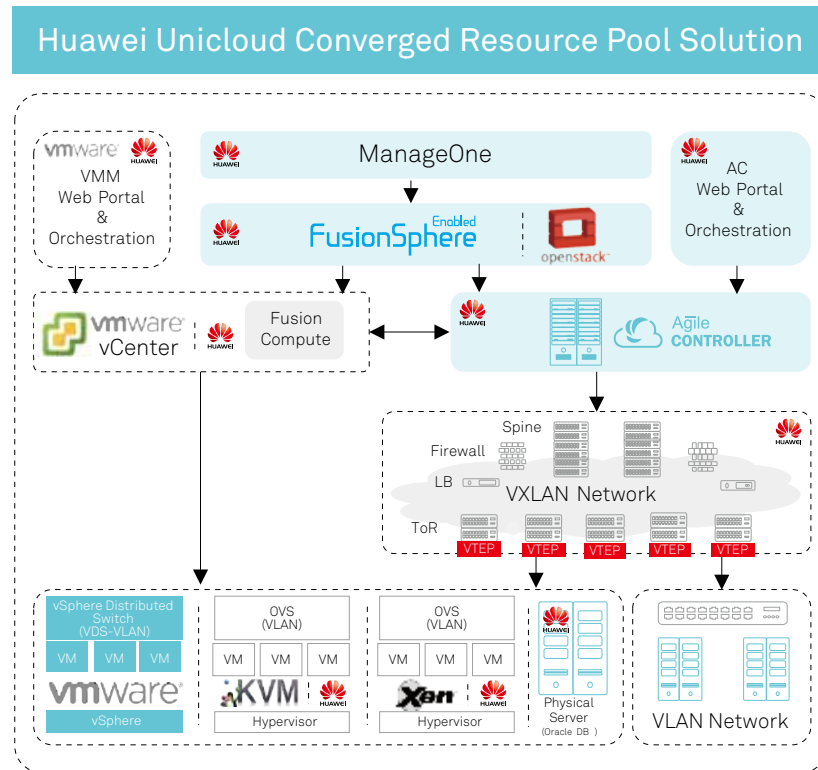
Huawei Cloud Fabric and VMware Converged Resource Pool Solutions



VMware Cloud Data Center & Huawei Joint Solution



Solution 1: Huawei Unicloud Converged Resource Pool Solution – Overview



Unified management of a converged resource pool

- Unified management of virtual and physical resources
- Multi-vendor device management
- Unified monitoring

FusionSphere Cloud Platform

- Standard OpenStack interface
Use APIs originating from the OpenStack community; quickly synchronize with latest version released on the OpenStack community.
- Support for third-party systems and devices
Compatible with the OpenStack ecosystem; supports heterogeneous hypervisor platforms and hardware devices.
- Hardening for commercial use, OpenStack enhancement
Easy deployment, easy O&M, high reliability, security hardening, automation, scalability

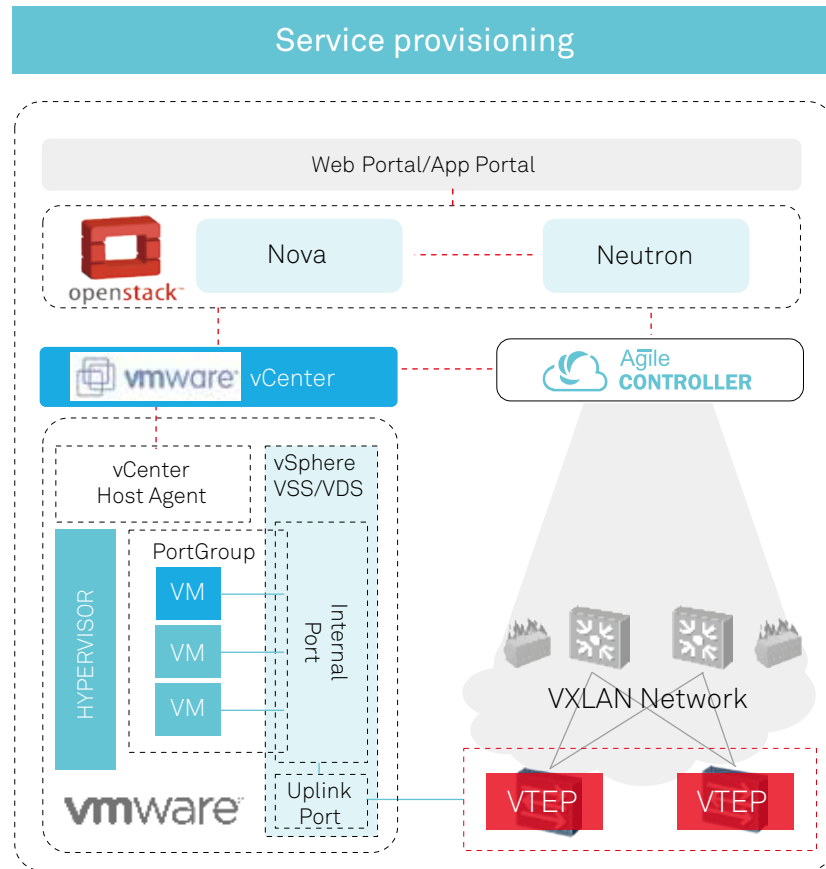
SDN-based hardware overlay network

- Automated network deployment
Open controller architecture, automated service orchestration and collaboration with VMM, high reliability and scalability of controller
- Elastic, reliable network
Ability to evolve to 40G/100G network, network virtualization, path optimization
- Refined network O&M
Visible: What you see is what you get; measurable packet loss ratio, latency, jitter; operability

Virtual and physical resource pool convergence

- Support for KVM, VMware, and Xen virtualization platforms
- Automated provisioning of bare metal servers
- Physical server connection with L2 bridges
- Automated resource pool management and resource sharing

Solution 1: Huawei Unicloud Converged Resource Pool Solution -- Service Provisioning



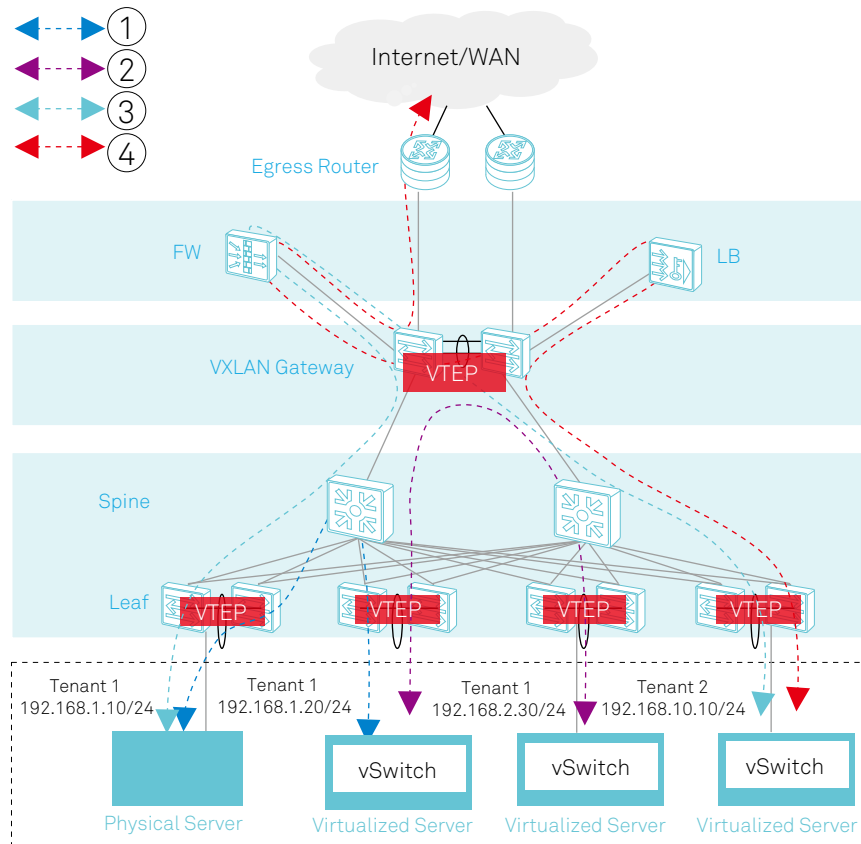
Step 1: Provision network resources.

- A Tenant Administrator creates a logical network on the service provisioning UI of the cloud platform.
- The Agile Controller interoperates with vCenter using Java interface to complete local network provisioning on VMware VDS.

Step 2: Provision compute resources.

- The Tenant Administrator provisions compute resources (bring VMs online/offline, migrate VMs), and binds VMs to the logical network on the cloud platform.
- Nova-Scheduler works with vCenterDriver to set up an independent resource pool with the vSphere cluster managed by vCenter. Nova-Scheduler does not select hosts in the vSphere cluster. vCenter schedules ESXi hosts based on dynamic entitlements and resource allocation settings, and provides LB and Dynamic Resource Scheduling (DRS) functions automatically.

Solution 1: Huawei Unicloud Converged Resource Pool Solution -- Traffic Model



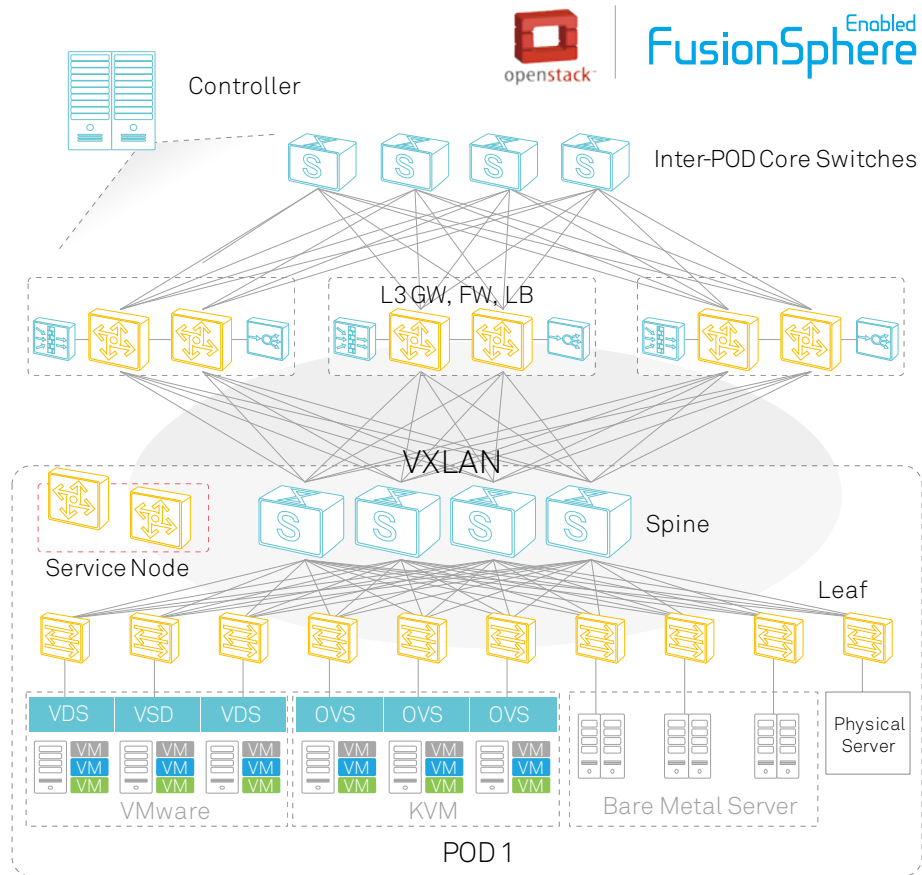
Types of traffic on a hardware VXLAN network

- Depending on traffic flow direction, traffic in a DCN is classified into east-west traffic (intra-DC communication) and north-south traffic (communication with external networks).
- Depending on transmission scope, traffic is classified into four types: 1. intra-subnet traffic within a tenant network; 2. inter-subnet traffic within a tenant network; 3. inter-tenant traffic; 4. traffic from outside DC.

Traffic transmission paths

- 1) Intra-subnet traffic within a tenant network is transmitted within a VXLAN L2 broadcast domain, and therefore is forwarded at L2.
- 2) Inter-subnet traffic within a tenant network is transmitted across VXLAN L2 broadcast domains. It needs to be forwarded by a VXLAN gateway but does not need to pass through a firewall.
- 3) Inter-tenant traffic needs to be forwarded by a VXLAN L3 gateway and processed by a firewall for secure isolation.
- 4) Traffic from outside DC needs to be forwarded to a tenant server through IPS/FW, LB, VXLAN gateway, and ToR switch.

Success Story: XX Bank of China's SDN Private Cloud



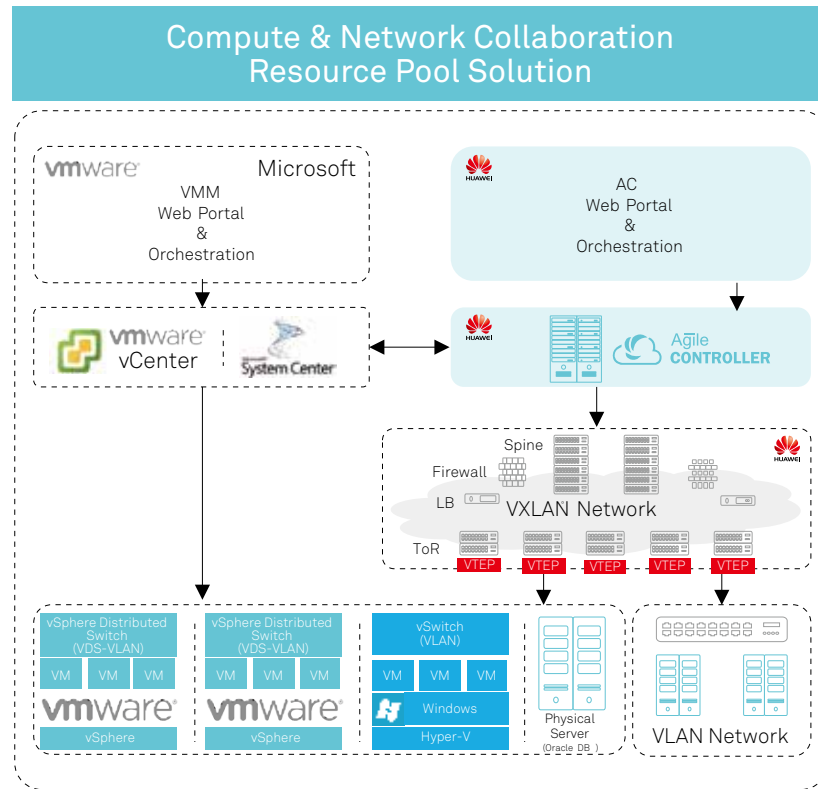
Solution

- The infrastructure uses a spine-leaf architecture, and the overlay network is a VXLAN-based large L2 network, with VTEPs deployed on leaf switches.
- The network uses centralized VXLAN L3 gateway deployment. More gateway groups can be deployed to support service expansion.
- Hardware FWs are used for north-south traffic, deployed in inline or bypass mode (inline routed mode recommended).
- The network controller works with an OpenStack cloud platform to automate network service deployment.

Benefits

- Hardware overlay networking ensures microsecond-level latency and jitter. Stacking and all-active gateway deployment provide high network performance and reliability.
- Centralized gateway deployment suits O&M models of traditional data center networks.
- Network resource sharing and automated service deployment improve the efficiency of application deployment.
- Multi-vendor virtualization platforms are converged to enable unified management of compute resource pool.

Solution 2: Compute & Network Collaboration Resource Pool Solution -- Overview



Centralized management of network resource pool

- Virtual and physical network resource polling, unified management, and on-demand allocation
- Application-based network service orchestration, drag-and-drop provisioning, what you see is what you get
- Unified network monitoring and management, application, logical, and physical networks mutually visible

Compute and network collaboration

- Compute and network resources cannot be managed by a unified cloud platform in a short period.
- Compute administrators create compute and storage resource pools on the VMM, without affecting running services.
- The Agile Controller creates a virtual network and is interoperable with multiple virtualization and physical host management platforms. It supports unified provisioning of physical and virtual networks and enables network policies to be migrated with VMs.

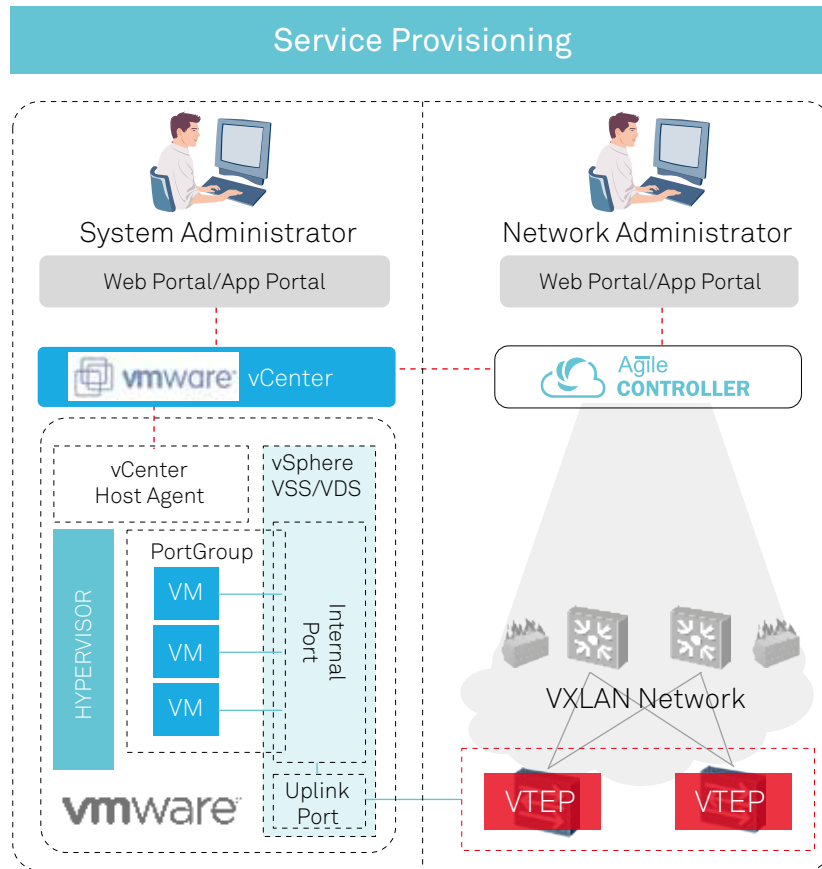
SDN-based hardware overlay network

- Automated network deployment
Open controller architecture, automated service orchestration and collaboration with VMM, high reliability and scalability of controller
- Elastic, reliable network
Ability to evolve to 40G/100G network, network virtualization, path optimization
- Refined network O&M
Visible: What you see is what you get; measurable packet loss ratio, latency, jitter; operability

Virtual and physical resource pool convergence

- Support for VMware and Hyper-V virtualization platforms.
- Physical server connection with L2 bridges
- Interoperable with traditional networks, allowing for future network convergence
- Automated resource pool management and resource sharing

Solution 2: Compute & Network Collaboration Resource Pool Solution -- Service Provisioning



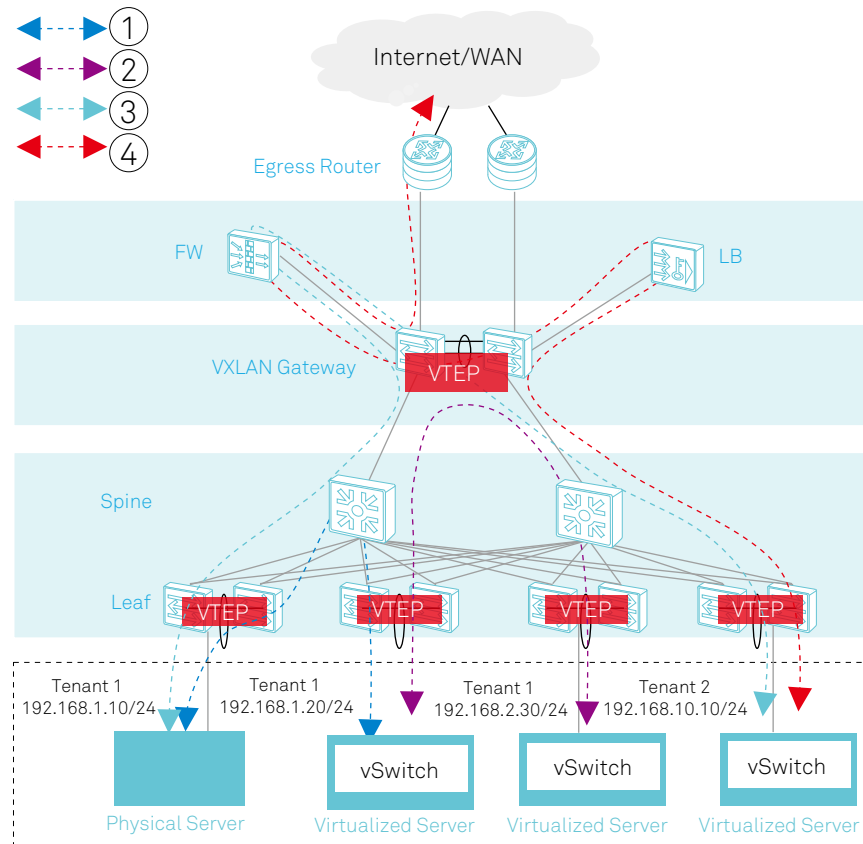
Step 1: Provision network resources.

- A Network Administrator creates a logical network on the network provisioning UI of the Agile Controller.
- The Agile Controller interoperates with vCenter using SOAP interface to complete local network provisioning on VMware VDS.

Step 2: Provision compute resources.

- A System Administrator provisions compute resources (bring VMs online/offline, migrate VMs), and binds VMs to the logical network on vCenter.
- The Agile Controller interoperates with vCenter using SOAP interface to detect VM login/logout/migration events and dynamically delivers configuration to network devices based on VM locations.
- The Network Administrator is unaware of the System Administrator's operations on VMs.

Solution 2: Compute & Network Collaboration Resource Pool Solution -- Traffic Model



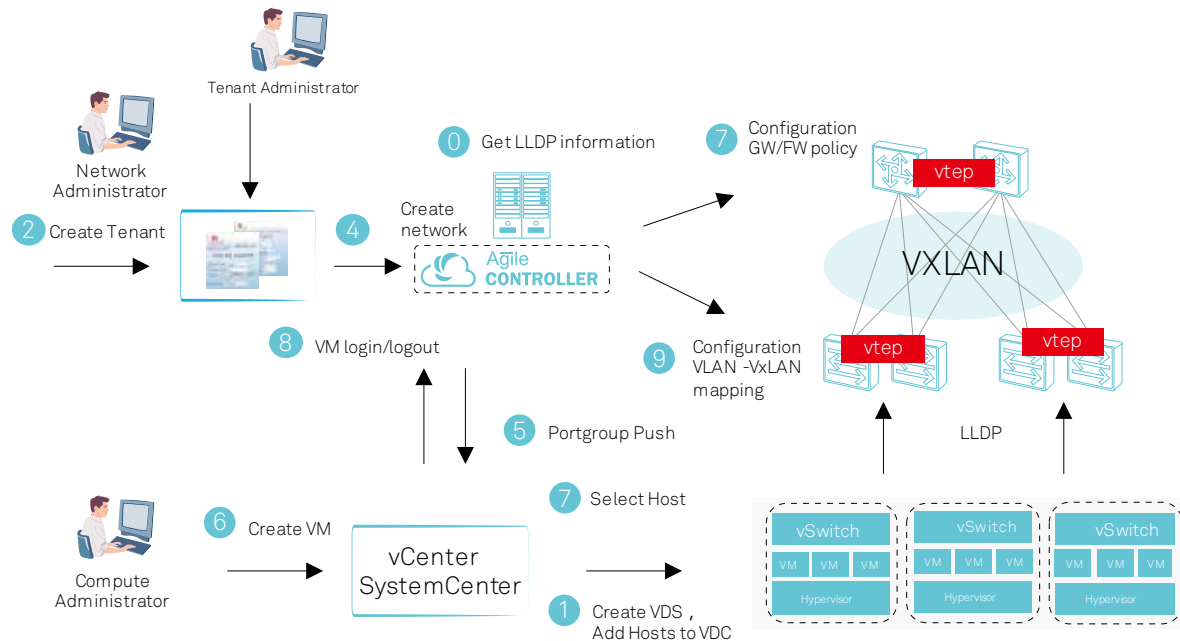
Types of traffic on a hardware VXLAN network

- Depending on traffic flow direction, traffic in a DCN is classified into east-west traffic (intra-DC communication) and north-south traffic (communication with external networks).
- Depending on transmission scope, traffic is classified into four types: 1. intra-subnet traffic within a tenant network; 2. inter-subnet traffic within a tenant network; 3. inter-tenant traffic; 4. traffic from outside DC.

Traffic transmission paths

- 1) Intra-subnet traffic within a tenant network is transmitted within a VXLAN L2 broadcast domain, and therefore is forwarded at L2.
- 2) Inter-subnet traffic within a tenant network is transmitted across VXLAN L2 broadcast domains. It needs to be forwarded by a VXLAN gateway but does not need to pass through a firewall.
- 3) Inter-tenant traffic needs to be forwarded by a VXLAN L3 gateway and processed by a firewall for secure isolation.
- 4) Traffic from outside DC needs to be forwarded to a tenant server through IPS/FW, LB, VXLAN gateway, and ToR switch.

Success Story: A Large Nordic ISP's DC SDN Solution



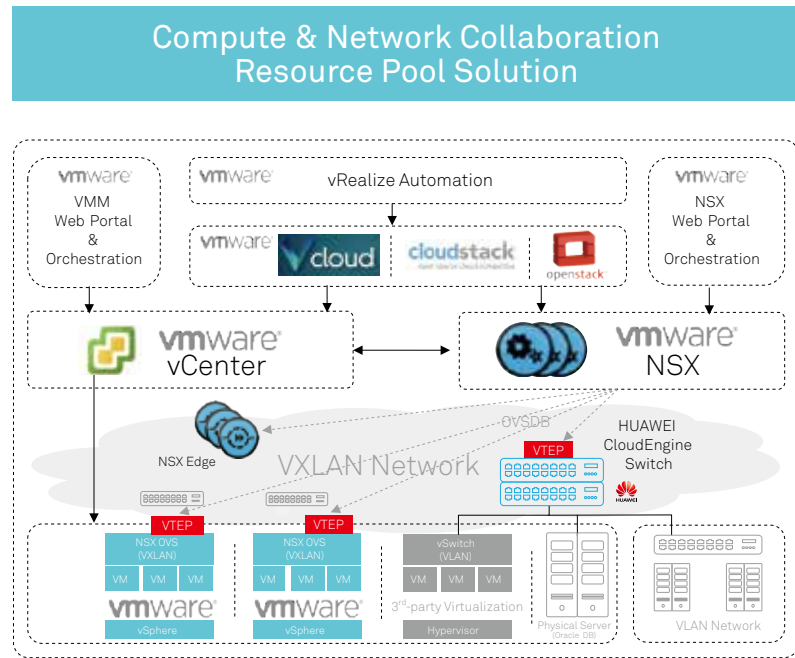
Solution

- The customer is the largest IT services provider in Northern Europe and provides multi-tenant and hosting services on a public cloud.
- The Agile Controller works with a cloud platform to complete service provisioning automatically.
- Tenants apply for cloud services on the self-service portal page, without intervention by Administrators.
- The fabric network is deployed using VXLAN technology.
- Agile Controller implements tenant isolation and controls tenants' access to branches and Internet.

Benefits

- Simplified O&M: Services of tenants can be provisioned automatically based on requirements of tenants.
- Openness and scalability: The service is compatible with the customer's service platform.
- Flexible deployment: As the customer's network and IT departments have not been converged, network automation can be implemented first, and network and compute resource convergence can be implemented later.
- Resource convergence: Physical servers and virtualization platforms VMware vSphere and Microsoft Hyper-V can be used to set up a compute resource pool.

Solution 3: VMware Cloud Data Center & Huawei Joint Solution – Overview



High-performance VXLAN GW

- Interconnection between VXLAN and non-VXLAN networks
- High-performance communication between servers



Flexible service deployment

- Flexible VM migration, automatic network resource scheduling for migrated VMs
- Unified VXLAN topology view

VMware vRealize cloud management platform

- Reuse the existing infrastructure, management tools, and management experience.
- Automatic deployment in various environments reduces development cost.

NSX-based software overlay network

VMware NSX is a transformative system that realizes the full potential of software-defined data centers by creating and operating a network on existing hardware devices.

- Agility and simplified O&M model

NSX can be deployed on the physical infrastructure without interrupting running services. Similar to server virtualization, NSX virtualizes a physical network into a transmission resource pool and uses the policy-driven model to provide network and security services.

- Advanced network connection and security service platform

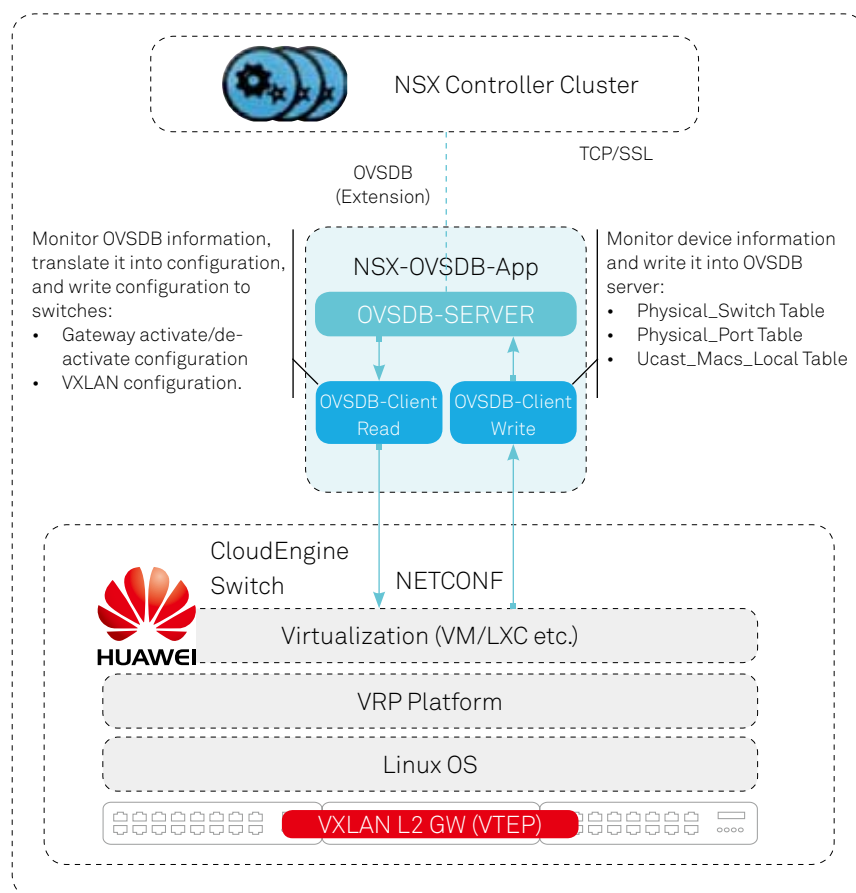
Traditional networks cannot change quickly to adapt to changes in data center services and are subject to configuration mistakes. NSX can configure virtual networks dynamically and automatically, enabling you to add virtual and physical services flexibly based on your service requirements.

Huawei and VMware Joint Solution

- CloudEngine data center switches use open software architecture and provide programmability using container technology.
- With a plug-in installed, CloudEngine data center switches can interoperate with NSX to enable communication between physical and virtual servers, and between VXLAN and traditional networks.
- CloudEngine switches ensure access layer reliability using stacking and M-LAG technologies.
- NSX centrally manages the overlay network and can connect to multi-vendor VAS devices to provide L4-L7 services.

Solution 3: VMware Cloud Data Center & Huawei Joint Solution— Interoperation

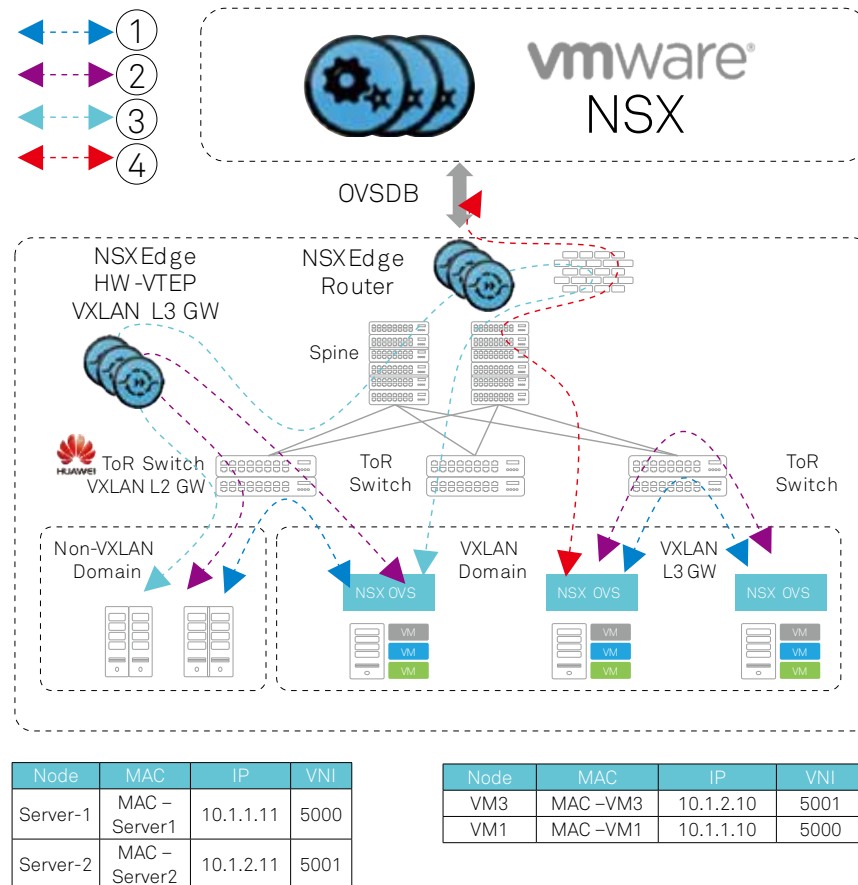
NSX and CloudEngine switch interoperation framework



Open architecture of Huawei CloudEngine switches



Solution 3: VMware Cloud Data Center & Huawei Joint Solution-- Traffic Model



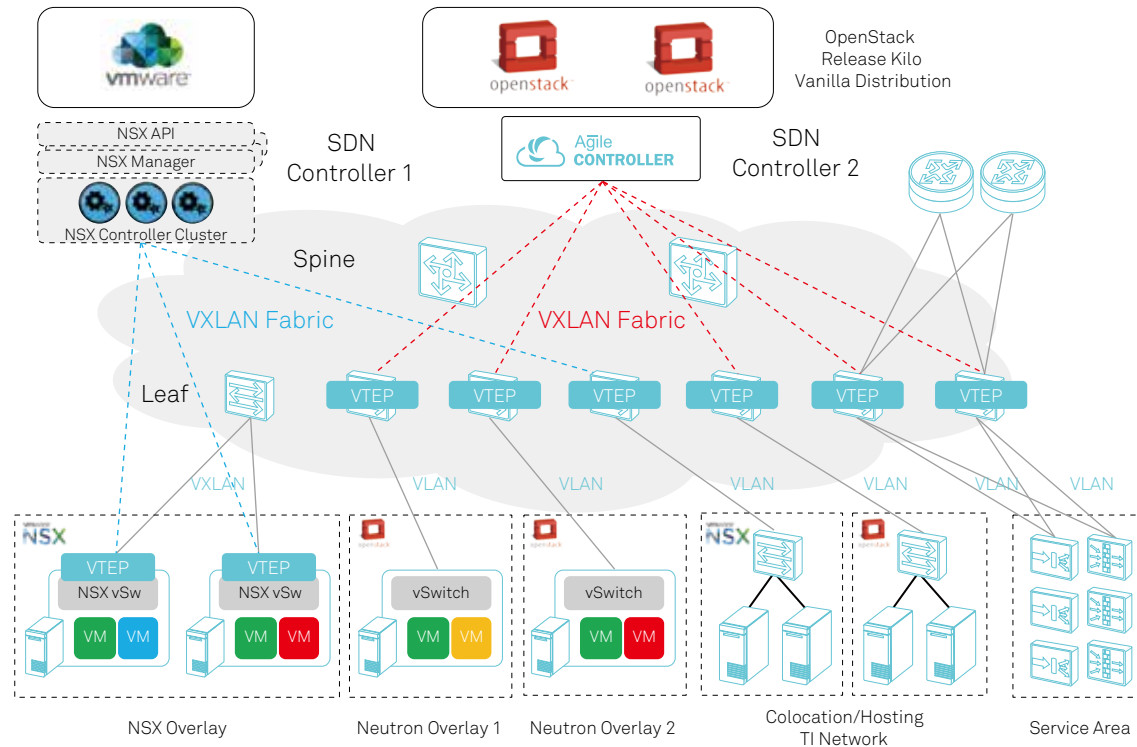
Types of traffic on a hybrid VXLAN network

- Depending on traffic flow direction, traffic in a DCN is classified into east-west traffic (intra-DC communication) and north-south traffic (communication with external networks).
- Depending on transmission scope, traffic is classified into four types: 1. intra-subnet traffic within a tenant network; 2. inter-subnet traffic within a tenant network; 3. inter-tenant traffic; 4. traffic from outside DC.

Traffic transmission paths

- 1) Intra-subnet traffic within a tenant network is transmitted within a VXLAN L2 broadcast domain, and therefore is forwarded at L2.
- 2) Inter-subnet traffic within a tenant network is transmitted across VXLAN L2 broadcast domains. It needs to be forwarded by a VXLAN gateway but does not need to pass through a firewall. NSX OVS provides VXLAN L3 gateway function and allows VMs on different hosts to communicate directly.
- 3) Inter-tenant traffic needs to be forwarded by a VXLAN L3 gateway and processed by a firewall for secure isolation.
- 4) Traffic from outside DC needs to be forwarded to a tenant server through IPS/FW, LB, VXLAN gateway, and VTEPs.

Success Story: XX Carrier's Public Cloud Project



Solution

- The customer is a Tier-1 carrier in Europe and built multiple data centers over the same physical network architecture to provide public and rack renting/hosting services.
- In the NSX overlay area, the NSX controller manages CloudEngine switches and enables communication between traditional hosting services and VM services.
- The OpenStack overlay area uses Huawei's hardware overlay SDN solution, in which the Agile Controller implements automatic service provisioning.

Benefits

- The VMware NSX and multiple OpenStack cloud platforms are deployed on the same fabric network. Both hardware and software overlay networks are deployed.
- The NSX-based software overlay solution provides unified management of virtual and physical networks. NSX manages PaloAlto and Fortinet firewalls to provide high-level security protection (micro-segmentation) for east-west traffic.
- The hardware overlay solution incorporates OpenStack and hosting services, and implements automated provisioning of physical and virtual networks.

Copyright © Huawei Technologies Co., Ltd. 2016. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademark Notice



HUAWEI, and  are trademarks or registered trademarks of Huawei Technologies Co., Ltd.

Other trademarks, product, service and company names mentioned are the property of their respective owners.

General Disclaimer

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.

HUAWEI TECHNOLOGIES CO.,LTD.
Huawei Industrial Base
Bantian Longgang
Shenzhen 518129,P.R.China
Tel: +86 755 28780808

www.huawei.com