

SDN as a Next Generation Open Network Infrastructure

A discussion of the technical requirements for realizing SDN technology, which enables flexible configuration and control of virtual networks via software, and a look at its future possibilities, as well as the features of the Stratosphere SDN Platform produced by Stratosphere Inc.

3.1 Introduction

SDN (Software Defined Networking)*1 is a concept for allowing users to control networks as they want via software programs. The IJ Group established Stratosphere Inc. in April of this year in a joint venture co-funded by ACCESS CO., LTD., and is working on developing infrastructure software for realizing SDN. In this report we provide a basic overview of the technical requirements associated with SDN, and look at the possibilities for developing next-generation cloud systems by meeting these requirements.

3.2. Demand for Virtualization in the Cloud Environment and Issues in Network Technologies

3.2.1 Demand for Cloud Services and the Limitations in Conventional Network Technologies

Users of cloud services and the operators who provide these services both expect cloud systems to meet the following requirements.

- Flexible, expandable system architecture that can be implemented swiftly
- Improved overall system utilization
- Standardized configuration and operating framework architecture
- High system reliability

Meanwhile, the VLAN (Virtual LAN) network virtualization technology used in current cloud systems has the following limitations.

- There are only 4,094 possible VLAN IDs
- It is not possible to create L2 segments between server clusters connected via router

Currently, local networks within server clusters are built with L2 switches, and are connected to other parts of the network via L3 routers. Within a cluster, a virtual L2 segment is configured for each user using VLAN. However, due to the above limitations, the number of L2 segments configurable within each cluster in this environment is restricted to a maximum of 4,094 VLAN IDs. Clusters that have reached the maximum number of VLANs are no longer able to accommodate new users, making it difficult to increase server utilization within the cluster. It is also not possible to extend the L2 segment of a given user beyond the cluster, leading to cases in which it is not possible to deal with requests to increase the number of VM (virtual machines) in an L2 segment.

Cloud operators face systems with limitations such as these when building multiple clusters with different policies and configurations, such as the number of servers per cluster, the number of VMs per server, or the specs of VMs, based on their quality of service definitions. This may lead to an overly complicated operating framework, increased costs, and lower overall reliability due to multiple different flows of operations.

3.2.2 From Network Virtualization to SDN

It is necessary to overcome the limitations of conventional network technologies for cloud systems to meet the new requirements arising from the popularization of cloud services. Therefore, new network virtualization technology that can replace conventional VLAN is essential for the evolution of cloud systems.

*1 Open Networking Foundation, "Software-Defined Networking: The New Norm for Networks" ONF White Paper, Apr. 2012

Based on the discussion in 3.2.1, it is clear that this new network virtualization technology must enable L2 segments to expand beyond physical network boundaries, and overcome the constraints of VLAN to accommodate larger numbers of user networks.

It is necessary for cloud computing technology and future cloud systems to provide continuous stable operation while carrying out the construction and management of large quantities of servers, vast numbers of VM running on these, and the complex networks connecting them. This means a considerable portion of system and network management must be automated, and operation must be made scalable. SDN is a concept devised to meet these requirements. SDN is a software-based solution that enables the construction of flexible virtual networks not subject to the limitations of conventional VLAN and L2 boundaries, as well as the automation of configuration management and operation for networks as a whole, including virtual networks. By defining APIs that enable these functions to be used from other systems, and providing them to users, it also allows users to utilize network infrastructure flexibly while constructing new services on cloud infrastructure for themselves (Figure 1).

We summarize the technologies necessary for realizing SDN below.

1. Virtualization technology that can operate a large number of logical networks independently of each other on a single physical network
2. Technology for managing the accommodation and configuration of virtual networks on a physical network, network object models for managing physical network resources shared among virtual networks, and APIs
3. Technology for the centralized remote control of network device settings and operation, to allow overall operation of multiple logical networks implemented on a physical network

With regard to 1 and 3, there is the OpenFlow*² protocol, which originated from the Clean Slate Program, an interdisciplinary research program at Stanford University. OpenFlow enables network virtualization by implementing a new kind of packet forwarding behavior. It does this by remotely setting rules known as flow rules that stipulate packet forwarding behavior for each flow in detail to network devices.

Meanwhile, new protocols for implementing virtual L2 networks over IP networks such as VXLAN*³, NVGRE*⁴, and STT*⁵ have also been proposed. These protocols establish IP tunnels between servers, enabling VM on the servers to connect to the same virtual L2 network. VLAN constraints are overcome by implementing new, larger ID fields such as VNI (VXLAN Network Identifier) or VSID (Virtual Subnet ID) to differentiate between logical networks.

There are two methods for achieving network virtualization (Figure 2). The first is the Hop-by-Hop method, which involves setting flow rules for each network device to create virtual networks for each flow. There are also Edge Overlay methods such as VXLAN and NVGRE, which create virtual networks using IP tunneling.




Figure 1: Conceptual Diagram of SDN




Figure 2: Hop-by-Hop Method and Edge Overlay Method

*² Nick MacKeown, et al. "OpenFlow: Enabling Innovation in Campus Networks", Mar. 2008

*³ M. Mahalingam, et al. "VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", Internet-Drafts, Aug. 2012

*⁴ M. Sridharan, et al. "NVGRE: Network Virtualization using Generic Routing Encapsulation", Internet-Drafts, Jul. 2012

*⁵ B. Davie, et al. "A Stateless Transport Tunneling Protocol for Network Virtualization (STT)", Internet-Drafts, Sep. 2012

The Hop-by-Hop method enables detailed flow-based control of traffic over each network link, but flow rules must be set to all network nodes for consistent overall packet forwarding.

Meanwhile, with the Edge Overlay method, as long as users assigned a VM on each server and their corresponding VNI or VSID are managed, virtual networks can be created flexibly using IP tunnels. However, detailed control of packet forwarding over IP networks between servers is not possible. The Edge Overlay method is generally considered to be easier to implement, because existing IP networks can be used without modification. In contrast, the Hop-by-Hop method requires all network devices to be replaced with those that support OpenFlow control. It is also possible to install a software switch on the IP tunnel terminal server when using the Edge Overlay method, and control software switch operation using OpenFlow. It would be best to consider utilizing a hybrid of the Hop-by-Hop and Edge Overlay methods, combining the features of each approach.

3.3 Stratosphere SDN Platform (SSP)

Here we will examine the features of Stratosphere Inc.'s product.

We adopted the Edge Overlay method for implementing virtual networks. We initially adopted the easy-to-implement Edge Overlay method, and in the future we will work on function enhancements for supporting the Hop-by-Hop method, interconnection with backbone protocols such as MPLS (Multi-Protocol Label Switching) and PBB (Provider Backbone Bridge), and extension to WAN using SDN, to achieve comprehensive wide-area network virtualization.

SSP supports VXLAN and STT for its overlay protocols. The VXLAN VNI is a 24-bit identifier, so approximately 16 million virtual L2 segments can be supported.

SSP provides a SDN solution that enables seamless migration from existing environments by blending new technology with a variety of conventional network technology. Figure 3 illustrates an Edge Overlay configuration using VXLAN on SSP. The Compute Nodes are virtualized servers. Each server has a virtual switch installed in its hypervisor, and VXLAN tunnels are established between these virtual switches. During this process L2 segment identification is carried out on the virtual switches using VLAN. When packets are passed from the virtual switch to the VXLAN tunnel, the VLAN ID is converted to a VNI. This configuration makes it possible to reduce the number of flow rules maintained on the Compute Node significantly, minimizing the packet processing overhead caused by OpenFlow. The SSP approach does not allow more than 4,094 VLANs to be set in a Compute Node, but in practical use 4,094 VLANs would conceivably last for quite some time. Considering the fact that most systems are currently run using VLAN, opting for this configuration is likely to create less stress during implementation.




Figure 3: Conceptual Diagram of SSP Tunnel Configuration (Built-In Mode)




Figure 4: Separate Mode

SSP also features a modular component configuration (Figure 4), so an external switch running on a Compute Node's hypervisor (built-in mode) can be run on another server as necessary (separate mode), reducing the processing load of the virtual switch on the Compute Node. In some situations it is also possible to carry out overlay processing via hardware using hardware appliances.

3.4 Network Modeling and APIs

Abstract models of the logical and physical networks handled using SDN are needed to realize the technologies listed under 2 in section 3.2.2. By defining the logical networks used by cloud users and the physical network that accommodates them using these models, it is possible to handle the optimal arrangement of logical networks within a physical network, and control the sharing of resources on an SDN system^{*6*7}. Additionally, by defining methods for manipulating these network models as APIs, functions for configuring, managing, and controlling virtual networks can be provided to external applications as services.




Figure 5: Opening Up Cloud Infrastructure to Users

For SSP, three layers of APIs are defined (Figure 5). The top layer is for controlling elements such as the virtual L2 segments connecting the VM of end users, and the bottom layer is for controlling physical hosts and network devices for infrastructure operators. The middle layer, which is for segmenting pieces of infrastructure and accommodating multiple user level virtual networks, borrows chunks of resources from the infrastructure operator, and carries out the various operations required by service providers, such as reselling to end users. The provider-oriented APIs from the middle layer release pieces of infrastructure to users. Provider users utilize these APIs to construct virtual network services tailored to their respective purposes, and provide them to end users.

3.5 Conclusion

SDN infrastructure and APIs allow service providers and end users to enjoy access to abstracted resources on cloud infrastructure through APIs without having to worry about details such as where or how logical systems are physically located. Meanwhile, infrastructure operators can use infrastructure APIs to identify where and how virtual networks are run by service providers and end users on their infrastructure, and reallocate or optimize resources as necessary. Developing cloud infrastructure that realizes this kind of environment paves the way for next-generation open network systems that maximize cloud infrastructure utilization while also offering service providers and end users an unprecedented degree of freedom.

Author:



Toshiya Asaba

President and CEO, IJ Innovation Institute Inc. President and CEO, Stratosphere Inc. Mr. Asaba joined IJ in its inaugural year of 1992, becoming involved in backbone construction, route control, and interconnectivity with domestic and foreign ISPs. He was named IJ director in 1999, and executive vice president in charge of technical development in 2004. When the IJ Innovation Institute Inc. was founded in June 2008, Mr. Asaba became its president and CEO. When Stratosphere Inc. was founded in April 2012, he also became president and CEO of that organization.

*6 T. Koponen, et al., "Onix: A Distributed Control Platform for Large-scale Production Networks" Operating Systems Design and Implementation (OSDI), Oct. 2010

*7 Asaba, "Possibilities of SDN technologies and vision for the next generation cloud" EIC Communication Society / Technical Committee on Internet Architecture, Sep. 2012