Secure Data Center Networking with Open vSwitch

Pa Pa Hlaing
University of Computer Studies, Yangon
ms.papahlaing@gmail.com

Abstract

Data centers have become the next-generation computing platforms for enterprises and Internet users. This is primarily due to the economic and technical advantages of resource sharing in data centers. By sharing computing and storage resources through services such as cloud computing or software-as-a-service (SaaS), users can amortize the cost of hardware and software. Because of the virtualization of resources, a new virtualized network access layer has been introduced to interconnect VMs within the data centers. In data center, hosts have been recently employed virtual switch to interconnect virtual machines (VMs) within data center networks. Virtual Switch is essential to control and manage VM within the hosts. Open vSwitch, a network switch specifically built for Xen virtualization environment is presented. The design and advantages of Open vSwitch is described. And then secure VM networking is proposed by using Open vSwitch. By combining Intrusion Detection System with Open vSwitch, more secure virtual networking can be established because Open vSwitch can get inter VM traffic logs directly and can perform more controls VM communication.

1. Introduction

Network virtualization in data centers has recently drawn much attention from the research community as its requirements and opportunities are far different from previously studied areas. The virtualization of the network layer in data center networks (DCNs) is expected to support agility and isolation of VMs. Most physical switches deployed in conventional DCNs are not designed for either supporting such unique VM requirements or flexible enough to augment new functionalities [7]. At the same time, VM networking has inherently unique characteristics, such as the awareness of the migration of VMs and their multicast membership. Therefore, much research has begun exploring the opportunity to introduce a new, flexible and programmable networking layer based on the knowledge of VMs.

Modern data center networks consist of both physical networks connected by switches and virtual networks formed by VMs running inside physical hosts. Inside one computer, many VMs can exist (as many as 120 VMs per host), each of which has at least one virtual network interface card (VNIC). The VNICs communicate with external networks through the host's physical NICs. The traffic multiplexing between the VNICs and physical NICs is achieved with a software layer in the host. This layer of software can be either Ethernet bridges or a virtual switch [6]. Figure 1 illustrates the architecture of a virtual switch residing in a host.

A virtual switch inside a host consists of fast-path and slow-path components, similar to a physical switch [2]. The fast path includes typical packet processing in a physical switch, including virtual Local Area Network (VLAN) packet encapsulation, traffic statistics collection, quality-of-service enforcement, and packet forwarding based on forwarding tables. The slow path is designed to support switch configuration and control.

Virtual switch can greatly complicate the ability to manage and monitor networks and applications. Managing VMs can very often be

done from the network. By using virtual switch, they increase visibility into inter-VM traffic through standard methods such as NetFlow and mirroring. Virtual switches also implement traffic policies to enforce security and quality-of-service requirements.

Virtual switch can greatly complicate the ability to manage and monitor networks and applications. Managing VMs can very often be done from the network. By using virtual switch, they increase visibility into inter-VM traffic through standard methods such as NetFlow and mirroring. Virtual switches also implement traffic policies to enforce security and quality-of-service requirements.

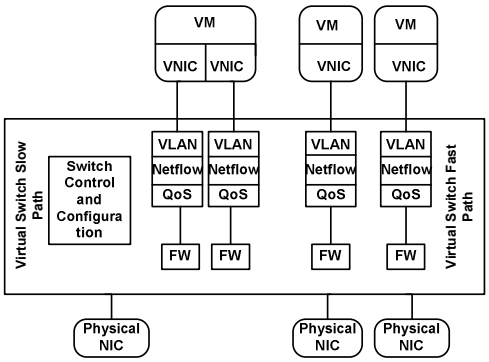

Figure 1. The architecture of a virtual switch residing in a host

Open vSwitch is a representative software virtual switch and it is a production quality, multilayer virtual switch licensed under the open source Apache 2.0 license [9]. It is designed to enable massive network automation through programmatic extension, while still supporting standard management interfaces.

Many designers have been proposed virtual switches including VMware's vNetwork Distributed Switch [12] and Cisco Nexus v1000 [3]. These designs implement in-host software based virtual switches inside either OS kernels or the hypervisors of VMs. These virtual switches will be discussed in next section.

The rest of the paper is organized as follows. The next section, section 2 reviews related work. Virtual switch, Open vSwitch, is described in Section 3. In Section 4, secure VM networking, relying on the proposed Open vSwitch is described. Finally, Section 5 concludes the paper.

2. Related Work

All useful virtualization platforms have some form of networking support. For example, VMware's recently released distributed switch [12], and Linux-based virtualization environments generally use the bridging code or the Virtual Distributed Ethernet (VDE) switch [11]. Networking vendors are also starting to create virtual switches [3] and other virtual network devices [1] for the virtualization layer. In most cases, these approaches use standard L2/L3 forwarding and standard interfaces for management, and have not demonstrated the flexibility needed to overcome some of the challenges.

Some designers note that the virtual switch shares and competes for the host's resources (CPU and memory) with the VMs running on the same host. It's unknown how virtual switch performance scales. As the first step of offloading virtual switching, a prototype of an OpenFlow switch has been implemented using a network processor (NP)-based NIC [4]. In this design, the OpenFlow flow table resides in the NP-based NIC, instead of the host memory.

In [7], Yan Luo proposes to offload the virtual switching onto the programmable NICs (PNICs) to achieve scalable VM networking. PNIC is defined as a NIC with programmable processors and memory units such as Netronome NFE-i8000 [8] and Net FPGA [5]. A PNIC often resides in a physical host as a powerful NIC to communicate with external networks. With PNICs, the virtual switch does not have to rely on the host CPU and memory to multiplex packets. Instead, the virtual switch can leverage the resources at the NICs, which can be seen as powerful line cards. PNIC is commercial hardware device.

This paper proposes secure VM traffics by using virtual switch, Open vSwitch. Open vSwitch is open source and compatible with Xen infrastructure. This design implement in-host

software-based virtual switch inside OS kernel of hosts [10]. In doing so, they can available control VM' traffic on the hosts, instead of relying solely on the dedicated physical switches. While the in-host virtual switching adds additional workloads to the host CPUs, this kind of approach takes advantages of the awareness of VM activities and host events.

3. Open vSwitch

3.1. Overview

Open vSwitch is a multilayer virtual switch designed with flexibility and portability in mind. It supports the features required of an advanced edge switch: visibility into the flow of traffic through NetFlow, sFlow, and mirroring (SPAN and RSPAN); fine-grained ACLs (Access Control Lists) and QoS (Quality-of-Service) policies and support for centralized control. It also provides port bonding, GRE and IPsec tunneling, and per-VM traffic policing. Since then, a configuration database has been introduced, which allows for the use of multiple simultaneous front-ends.

Open vSwitch is backwards-compatible with the Linux bridge, which many Linux-based hypervisors use for their edge switch. This allows it to be a drop-in replacement in many virtual environments. It is currently being used in deployments based on Xen, XenServer, KVM, and VirtualBox. The Xen Cloud Platform and upcoming versions of XenServer will ship with Open vSwitch as the default switch. Open vSwitch is also being supported to non-Linux hypervisors and hardware switches, due to its commercial friendly license, modular design, and increasing feature set.

3.2. Benefits of Open vSwitch

The benefits introduced by Open vSwitch virtual switching are as follows:

Cost

Deployment of virtual switching has minimal hardware cost. Virtual switching such as Open vSwitch is primarily a software feature. Deploying it requires installing software in the hypervisor layer and upgrades require only a software update.

Performance

A virtual switch's greatest strength is in VM-to-VM traffic, which never needs to hit the hardware wire. The communication path is defined by the Open vSwitch, with its high bandwidth, low latency, and negligible error rate, rather than network capabilities. Throughput is higher, errors due to corruption, congestion is unnecessary.

Visibility

Traditional switches provide network administrators very little insight into the traffic that flows through them. An Open vSwitch directly interacts with all virtual machines and can gather any data it needs without intermediate layers or inferences. The Open vSwitch's visibility into the source of virtual network traffic gives it the ability to affect traffic before it enters the network.

Control

Using the visibility advantages of the Open vSwitch, these switches are in a prime position to enforce network policy. The available rich sources of information give network administrators the ability to make fine-grained rules for handling traffic. In addition, by applying policy at the Open vSwitch, it is possible to drop unwanted traffic as soon as possible.

3.1. Integration with XenServer

Open vSwitch works seamlessly with XenServer. Future versions of XenServer will ship with Open vSwitch as the default. It is also fully compatible with XenServer 5.6, the currently shipping version. XenServer is built around a programmatic interface known as XAPI (XenServer API). XAPI is responsible for managing all aspects of a XenServer, including VMs, storage, pools, and networking. XAPI

provides an external interface for configuring the system as well as a plug-in architecture that allows applications to be built around XenServer events. XAPI notifies Open vSwitch of events that may be of interest. The most important of these are related to network configuration. Internal and external networks may be created at any time. External networks support additional configuration such as port bonding and VLANs. These networks are essentially learning domains for which virtual interfaces may be attached.

XAPI notifies Open vSwitch when bridges should be created and interfaces should be attached to them. When it receives such a notification, Open vSwitch queries XAPI to determine greater context about what is being requested. For example, when a virtual interface is attached to a bridge, it determines the VM associated with it. Open vSwitch stores this information in its configuration database, which notifies any remote listeners, such as a central controller.

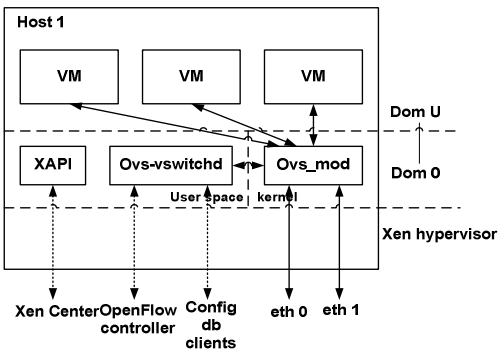

Figure 2. Open vSwitch on XenServer

Figure 2 shows Open vSwitch integration with Xen server. Each unprivileged virtual machine (DomU) has one or more virtual network interfaces (VIFs). VIFs communicate with the Open vSwitch “fast path” kernel module ovs_mod running in the control VM (Dom0). The kernel module depends on the userspace ovs-switchd “slow path”, which also implements the OpenFlow and configuration database protocols. XAPI, the XenServer management stack, also runs in Dom0 userspace.

4. Secure VM networking

Secure VM networking is important for virtualized Data Center. Virtualization refers to the abstraction of logical resources away from their underlying physical resources in order to improve agility and flexibility, reduce costs and thus enhance business value. In a virtualized environment, computing can be dynamically created, expanded, shrunk or moved as demand varies.

Virtualized Data Center require significant amounts of computing power, especially as those virtual spaces become large or as more and more users log in. Massively multiplayer online games are a good example. Several commercial Data Center have as many as nine million registered users and hundreds and thousands of servers supporting these environments.

In these Virtualized environments, Intrusion Detection System is essential for VM networking. For example, network attack such as DoS (Denial of Service) attacks reduce the bandwidth and increase the congestion causing poor service to the need. Massive distributed DoS attacks have the potential to severely decrease backbone availability and can virtually detach a network from the internet.

To detect and prevent attack, Intrusion Detection system must log the network traffics. Commercial products for VM networking such as commercial physical switch provide traffic in-bound and out-bound for managing. However, in virtualized environment, virtual switch is more suitable than physical switch to monitor and control traffic. Virtual switch like Open vSwitch has more visible inter-VM traffic and more cost effective to detect and prevent attack.

4.1 Intrusion Detection System

Intrusion detection systems attempt to detect and report whether a host has been compromised by monitoring the host’s observable properties, such as internal state, state transitions (events), and IO activity.

A host-based intrusion detection system offers a high degree of visibility as it is

integrated into the host it is monitoring, either as an application, or as part of the OS. The excellent visibility afforded by host-based architectures has led to the development of a variety of effective techniques for detecting the influence of an attacker, from complex system call trace analysis.

Network-based intrusion detection systems offer significantly poorer visibility. They cannot monitor internal host state or events, all the information they have must be gleaned from network traffic to and from the host. Limited visibility gives the attacker more room to maneuver outside the view of the IDS. An attacker can also purposefully craft their network traffic to make it difficult or impossible to infer its impact on a host.

If the IDS reside on the host, it has an excellent view of what is happening in that host's software, but is highly susceptible to attack. On the other hand, if the IDS reside in the network, it is happening inside the host, making it more susceptible to evasion. IDS with Open vSwitch retain the visibility of host-based IDS. But pull the IDS outside of the host for greater attack resistance. Using this approach allows the system to isolate the IDS from the monitored host but still retain excellent visibility into the host's state.

4.2 Intrusion Detection System with Open vSwitch

Open vSwitch is installed on the Xen hypervisor to log the network traffic in-bound and out-bound into the database for auditing as shown in Figure 3. Open vSwitch can provide real-time traffic with great visibility within VMs. It has greatest strength in VM to VM traffic.

The traffic packets are examined in real-time by the intrusion detection system for a particular type of attack based on predefined rules. The rules are defined based on well known attack strategies by the intruders. The Intrusion Detection System could determine the nature of attack and is capable of notifying centralized management center the amount security risks involved.

Virtualization platforms typically include a centralized management system for managing the deployment of physical hosts and VMs, such as Xen Center. Xen Center can provide various mechanisms of grouping and maintaining metadata about managed VMs, hosts, storage repositories, and so on. Xen Center can also provide access to alert that are generated when noteworthy things happen and can perform emergency response to the attack. The Intrusion Detection System's security management function connects with this virtualization management system called Xen Center to control the hosts and virtual machines security.

Intrusion Detection System has more insight into using Open vSwitch. The closer the measurement point is to the component being measured, the richer the information available. Open vSwitch directly interacts with all virtual machines and can gather any data it needs without intermediate layers or inferences. VM networking with Open vSwitch can support secure VM connectivity in the Xen virtualization environment.

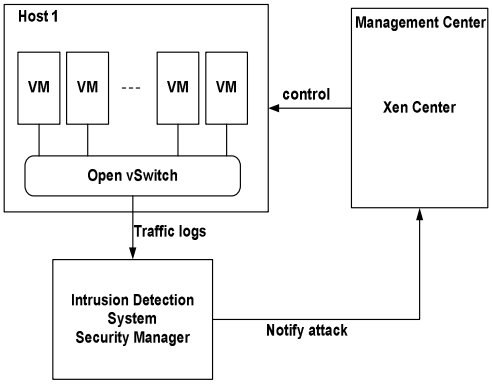

Figure 3. Architecture of IDS with Open vSwitch

5. Conclusion

CPU virtualization has evolved rapidly over the last decade, but virtual networks have lagged behind. The switches that have long shipped in

hypervisors were often little more than slightly modified learning bridges. They lacked even basic control and visibility features expected by network administrators. Advanced switches, such as Open vSwitch, answer many of their shortcomings and provide features previously only available in high-end hardware switches. By using Intrusion Detection System with Open vSwitch, the network administrator in Data Center can easily monitor and control VM networking and can detect and prevent attacks easily.

References

[1] Altor Networks. http://altornetworks.com, July 2009.

[2] B. Pfaff, J. Pettit, T.Koponen, K. Amidon, M. Casado, and S. Shenker. “Extending Networking into the Virtualization Layer”. In ACM HotNets, New York, NY. October 2009.

[3] Cisco. Cisco Nexus 1000V Series Switches. http://www.cisco.com/en/US/products/ps9902/, 2009.

[4] C. Pablo, L. Yan, M. Eric, and O. Julio. “Accelerating OpenFlow Switching with Network Processors”. In ACM ANCS’09 Workshop, August 2009.

[5] J.W. Lockwood, M. Nick, W. Greg, G. Glen, H. Paul, N. Jad, R. Ramanan, and L. Jianying. “Netfpga-an open platform for gigabit-rate network switching and routing”. In MSE’ 07: Proceedings of the 2007 IEEE International Conference on Microelectronic Systems Education, pages 160-161, San Diego, CA, USA, 2007. IEEE Computer Society.

[6] L. Yan. “Network I/O virtualization for cloud computing”. In IEEE Computer Society, October 2010.

[7] M. Eric, L. Yan, T.L. Ficarra. “Accelerate Virtual Switching with Programmable NICs or Scalable Data Center Networking”. In ACM VISA 2010, New Delhi, India, September, 2010.

[8] Netronome. Product Brief-NFE-i8000 Network Acceleration Card, 2006.http://www.netronome.com/.

[9] Nicira Networks, “Open vSwitch: An open virtual switch,” http://www.openvswitch.org/, May 2010.

[10] P. Justin, G. Jesse, B. Pfaff and C. Martin. “Virtual Switching in an Era of Advanced Edges”. Nicira Networks, Palo Alto, CA, 2010.

[11] R. Davoli. VDE: Virtual Distributed Ethernet. In Proc. Of TRIDENTCOM, Feb. 2005.

[12] VMWare vSphere: vNetwork Distributed Switch. http://www.vmware.com/products/vnetwork - distributed-switch/, 2010.