

DDoS Attack Protection in the Era of Cloud Computing and Software-Defined Networking

Bing Wang Yao Zheng Wenjing Lou Y. Thomas Hou
Virginia Polytechnic Institute and State University, Blacksburg, VA, USA
{bingwang,zhengyao,wjlou,thou}@vt.edu

Abstract—Cloud computing has become the real trend of enterprise IT service model that offers cost-effective and scalable processing. Meanwhile, Software-Defined Networking (SDN) is gaining popularity in enterprise networks for flexibility in network management service and reduced operational cost. There seems a trend for the two technologies to go hand-in-hand in providing an enterprise’s IT services. However, the new challenges brought by the marriage of cloud computing and SDN, particularly the implications on enterprise network security, have not been well understood. This paper sets to address this important problem. We start by examining the security impact, in particular, the impact on DDoS attack defense mechanisms, in an enterprise network where both technologies are adopted. We find that SDN technology can actually help enterprises to defend against DDoS attacks if the defense architecture is designed properly. To that end, we propose a DDoS attack mitigation architecture that integrates a highly programmable network monitoring to enable attack detection and a flexible control structure to allow fast and specific attack reaction. The simulation results show that our architecture can effectively and efficiently address the security challenges brought by the new network paradigm.

I. INTRODUCTION

As cloud computing provides on-demand, elastic, and accessible computing services, more and more enterprises begin to embrace this paradigm shift by moving their database and applications into the cloud. At the same time, another epochal concept of the Internet architecture comes to forefront, i.e., Software-Defined Networking (SDN). While cloud computing facilitates the management of computation and storage resources, SDN is proposed to address another laborious issue hindering the evolvement of today’s Internet, i.e., the complicated network management. Besides the fact that SDN has been proposed as a candidate of the next generation Internet architecture, companies like Google have already adopted SDN in their internal data centers. Thus, the arrival of the era when cloud computing and SDN go hand-in-hand in providing enterprise IT services is looming on the horizon.

Besides all the widely perceived benefits, the marriage between cloud computing and SDN may also introduce potential risks, especially on network security. Among all the network security problems, we first take a look at Denial-of-Service (DoS) attack. A DoS attack and its distributed version, Distributed Denial-of-Service (DDoS) attack, attempt to make a service unavailable to its intended users by draining the

system or network resource. Although network security experts have been devoting great efforts for decades to address this issue, DDoS attacks continue to grow in frequency and have more impact recently. Existing DDoS attack defense solutions (to list a few [1], [2], [3], [4]) assume a fully controlled network by the network administrators of enterprises. Therefore, the network administrators could place certain hardware pieces in the network to detect or mitigate DDoS attacks. However, in the new network paradigm of cloud computing and SDN, these assumptions no longer stand. Other researchers [5], [6] focus on exploiting the benefits of cloud or SDN to defend DDoS attacks. But their target victims still reside in the traditional network environment, which makes their solutions unsuitable for the new network paradigm. To the best of our knowledge, little effort in research community has been made to look into the potential problems or opportunities to defend DDoS attacks in the new enterprise network environment that adopts both cloud computing and SDN.

In this paper, we first analyze the impact of the combination of cloud computing and SDN on DDoS attack defense. We discuss the potential issues under this new paradigm as well as opportunities of defending DDoS attacks. Based on our analysis, we claim that if designed properly, SDN can actually be exploited to address the security challenges brought by cloud computing and the DDoS attack defense can be made more effective and efficient in the era of cloud computing and SDN. We then propose a new DDoS attack mitigation architecture using software-defined networking (abbreviated as DaMask) to demonstrate and substantiate our findings. Our contributions can be summarized as follows:

- 1) To the best of our knowledge, we are among the first to bring the attention of the impact on DDoS attack defense of the new network paradigm, which is a combination of cloud computing and SDN. Based on our analysis, we find that the marriage of SDN and cloud computing provides an unique opportunity to enhance the DDoS attack defense in an enterprise network environment.
- 2) To substantiate our claim, we propose DaMask, a highly scalable and flexible DDoS attack mitigation architecture that exploits SDN technique to address the new security challenges brought by cloud computing, in-

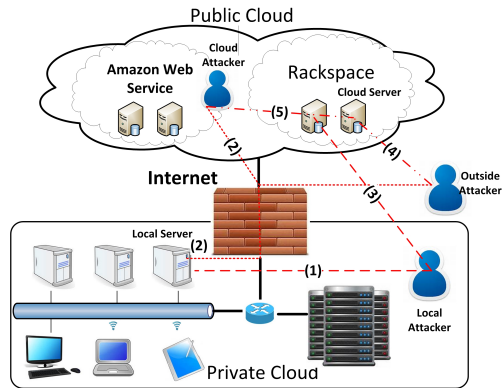


Fig. 1: The structure of a hybrid cloud, consisting of one private cloud and two public clouds. Five types of attack traffic are shown in the figure.

- cluding the extended defense perimeter and the dynamic network topological changes.
- 3) At last, we implement our proposed structure and performed a simulation based evaluation using the Amazon EC2 cloud service. The results show that our scheme works well under the new network paradigm and incurs limited computation and communication overhead, which is a crucial requirement of DDoS protection in cloud computing and SDN.

We organize the remainder of the paper as follows. We analyze the impact of cloud computing and SDN on DDoS attack defense in Section II. Based on our analysis, we formulate the problem and present our DaMask architecture design in Section III. Section IV presents the simulation setting and the results. Related work are reviewed and compared with our work in Section V. We draw concluding remarks in Section VI.

II. ANALYSIS

In this section, we briefly review cloud computing and SDN. Then we analyze the impact of the combined technologies on the network protection against DDoS attacks.

A. Cloud Computing

Cloud computing is a computing model which manages a pool of configurable computing resource. Cloud computing can be categorized as *public cloud*, *private cloud* and *hybrid cloud* in terms of deployment. While public cloud and private cloud are used by public and a single organization, respectively, hybrid cloud is a composition of public and private cloud infrastructures. As a result, hybrid cloud share the properties of both public cloud and private cloud. Hybrid cloud allows companies keeping their critical applications and data in private while outsourcing others to public. Thus, we focus on analyzing the impact of hybrid cloud on DDoS attack defense.

B. Impact of cloud computing on DDoS attack defense

Nowadays, attackers can launch various DDoS attacks including resource-focused ones (e.g. network

bandwidth, memory, and CPU) and application-focused ones (e.g. web applications, database service) from almost everywhere. To be realistic, we have to assume attackers can reside either in a private network, in a public network, or in both. To this end, we find the following properties of cloud computing affect DDoS attack defense.

- 1) Instead of users, cloud providers control network and computation resources, i.e., physical servers. This property differs from the system model in the traditional DDoS attack defense, where the protected application servers are within the defender controlled network.
- 2) Resource allocation and virtual machine migration are new sources of network topological changes from the defender's view. Moreover, the resource allocation and virtual machine migrations processes are fast-paced. The DDoS attack defense must be able to adapt to a dynamic network with frequent topological changes and still maintain high detection rate and prompt reaction capability.
- 3) All cloud users share the same network infrastructure of the cloud. This raises a reliable network separation requirement, which has not been considered in traditional DDoS attack defense. The enterprise must ensure its DDoS attack detection/defense operations neither affect nor be affected by other cloud users.

We illustrate these impacts using the example in Fig. 1. To ease the presentation, we denote an attacker in the private cloud of the enterprise network as a *local attacker*, an attacker in the off-site public cloud of the enterprise network as a *cloud attacker*, and other attackers as *outside attackers*. Similarly, we refer a server in the private cloud as a *local server* and a company's server in the public cloud as a *cloud server*. We consider two attacking scenarios. In the first attacking scenario, the victim server is within the private cloud. In the second one, the victim server resides in the public cloud.

In the first attack scenario, there are two types of attack traffic, i.e., (1) and (2) in Fig. 1. The enterprise's local DDoS attack defense system can detect the attack traffic (2), while the detection of the attack traffic (1) depends on whether the internal traffic is redirected to the DDoS attack defense system. Nevertheless, this scenario is similar to the traditional DDoS attack scenario. In what follows, we focus on two new challenges introduced in the second attacking scenario.

The first challenge is raised by the public accessibility of the cloud resources. We refer to this challenge as *extended defense perimeter*. There are three types of attack traffic, (3), (4) and (5). The enterprise's local defense can only examine and filter out the attack traffic (3) before the traffic leaves the local network. The defense offered by the cloud provider can check the attack traffic (4). However, more advanced attacks, such as the application-layer attacks which

target specific applications, can bypass the generic defense provided by the cloud. The most stealthy attack is type (5) because it is initialized from the same physical network or even the same physical machine on which the application is running. Most of these traffic is handled at local switches or hypervisors without going through the detection hardware.

The second challenge is raised by the rapid resource re-allocation. We refer to this challenge as *dynamic network topology*. This challenge makes the attack traffic (4) and (5) more difficult to handle because the enterprise's defense mechanism has to adjust to the network change caused by the physical location change of the virtual machine. The adaption must take effect in a short time period, for example, in milliseconds thanks to advances in live migration technology [7]. Moreover, because most of the topology changes are done by cloud provider without notifying the users, the DDoS attack defense mechanism needs to communicate with the cloud service provider to properly adapt the changes.

C. Software-Defined Networking

Unlike the well formatted data plane abstraction in the OSI model, the control plane of the Internet is composed of various complicated protocols for various network functions. Managing these protocols in a distributed manner becomes inefficient and error-prone. SDN is a network architecture that decouples the control plane and the data plane of network switches and moves the control plane to a centralized application called *network controller*. The network controller is in charge of the entire network through a vendor-independent interface such as OpenFlow [8], which defines the low-level packet forwarding behaviors in the data plane. Developers then can program the network from a higher level without concerning the lower level detail of packet processing and forwarding in physical devices.

D. Impact of SDN on DDoS attack defense

The most important two concepts of SDN are *control plane abstraction* and *network function virtualization*. They introduce following properties.

- *Centralized network control*: The centralized network operating system (NOS) connects to all the switches in the network directly. Thus, NOS can provide a global network topology along with the real-time network status.
- *Simplified packet forward*: The data plane in SDN simply forwards packets based on the forwarding policies generated by control programs.
- *Software based network function implementation*: Network functions originally implemented within a switch or a middle-box are implemented as control programs in SDN. These control programs reside above the NOS and communicate with switches remotely.

- *Virtualized network*: Similar to a hypervisor in hardware virtualization, the network virtualization hides the network topology from control programs so that network function developers can focus on the functionality implementation.

Implementing SDN affects the DDoS attack defense greatly in both directions. On the bright side, SDN makes advanced detection logic and rich subsequent processes easier to implement. On the downside, the devices or middle-boxes originally distributed within the network now need to be located above NOS. Compared with hardware-based packet processing, software processes packets is much slower. The network delay and traffic overhead caused by the communications between the control program, i.e., the DDoS attack defense schemes, and the switches, may become the new attack surface.

E. DDoS attack defense in cloud computing and SDN

Based on our analysis, cloud computing introduces new DDoS challenges, i.e., *extended defense perimeter* and *dynamic network topology* due to its new operation model. To effectively address these challenges, the cloud provider must be able to 1) easily delegate the control of its network to cloud users; 2) fast re-configure the control according to the network topology changes caused by dynamic allocations and migrations. On one side, we could benefit from the centralized network controller and the network virtualization of SDN. On the other side, SDN influences DDoS attack defense in negative ways as we discussed early. The negative impact of SDN mainly comes from the efficiency of processing packets using software, which may generate new attack surface and lead to single-point failure. When designing a DDoS attack defense solution in SDN, one must take the computation and communication overhead into the consideration so that no new security vulnerability is introduced. To sum up, we believe SDN technology will benefit the DDoS attack defense in cloud computing as long as the design could carefully handle the communication and computation overhead.

III. DAMASK DESIGN

A. Design Overview

Based on the analysis in Section II, we need to incorporate the DDoS attack defense into cloud computing and SDN. To successfully address the DDoS attack defense challenges in the new network environment, we must achieve the following objectives. First of all, the scheme must be *effective*. The design should be able to protect the services in both private and public clouds. It also should be able to adapt to the network topology changes and mitigate DDoS attacks efficiently. Secondly, the scheme should incur *small overhead*. The communication and computation overhead introduced by the architecture should also be limited to a small amount to be practical. Lastly, the deployment cost should be *inexpensive*. The solution should require as little deployment cost, such as additional hardware or changing existing protocols for both enterprises and cloud service providers, as possible.




Fig. 2: Workflow of DaMask.

To address the first challenge, our idea is to separate the enterprise’s network traffic from the main network by virtualizing the network. We call such a virtual network a *slice*. Then we let the cloud provider delegate the slice to the owner of this slice. Similar with the hardware or platform virtualization, a slice contains the network flows related to the enterprise only and is isolated from other slices. The strong isolation between different slices ensures that a slice is visible to its belonging company only. Therefore, operations performed on the slice are transparent to other cloud users.

For the second issue, we should select an efficient attack detection algorithm which involves as little information as possible to reduce the communication overhead. Meanwhile, the detection process itself must be fast enough to incorporate with the packet forwarding speed. Existing DDoS attack detection algorithms could serve the purpose as long as it does not depend on certain hardware. It is also worth mentioning that signature-based detection or anomaly-based detection or even a combined detection scheme can be used here.

To cope with the last issue, we need a rapid re-configuration scheme for each slice in the cloud. Given the nature of a virtualized slice, which is defined by its profile, our idea is to re-configure each slice profile when the virtual machine migration is taking place. Because the cloud provider virtualizes the network, he can track all the enterprises’ controllers, and re-configure the profile of a slice when a migration is about to happen. By applying the new slice profile, the cloud provider ensures the right enterprise getting the control of the proper slice.

B. Workflow of DaMask

To substantiate our previous claim, we propose a DDoS attack mitigation architecture, named DaMask. The DaMask architecture has three layers, network switches, network controller, and network applications. The main functions of the DaMask are DDoS detection and reaction. There are two separate modules in the DaMask, *DaMask-D*, a network attack detection system, and *DaMask-M*, an attack reaction module. We present the workflow of DaMask in Fig. 2.

1) *DaMask-D module:* The DaMask-D module is an anomaly-based attack detection system. We argue that although signature-based attack detection could also work, they are not efficient. The reason is that,

in SDN, the responsibility of generating a packet signature moves from a switch or a middle-box to a remote control program, which not only processes slower than hardware, but also requires all the packets to be redirect to it. Therefore, we focus on anomaly-based detection. Now we assume we already have a detection algorithm implemented (this can be done in an offline process as shown in Fig. 2).

In online phase, when a new packet arrives at the switch, the cloud provider first decides which slice the packet belongs to. Then the cloud provider notifies the corresponding NOS of the slice. After receiving the notification, the slice owner’s NOS checks whether the packet belongs to an existing flow¹. If so, it updates the flow statistic, otherwise it build a new flow record. Then we query the detection model with the updated or the new flow static. If the query result indicates an *attack*, DaMask-D issues an *alert* and forwards the alert along with the packet info to the DaMask-M module. If the query result is *normal*, the packet is forwarded to its intended destination. Occasionally, the detection model cannot determine the attack type of a packet if the packet belongs to a new type of attack. In that case, the packet needs to be further analyzed. The analysis result is then used to update the detection model through a *model updating* process.

2) *DaMask-M module:* The DaMask-M module is an attack reaction system. In the existing work of DDoS protection in today’s Internet, the reaction options are simple and limited, because advanced post-processing logic requires switches working together in a distributed manner. Implementing and managing such functions are time-consuming and error-prone. In SDN, we can implement those sophisticated logic such as quarantine of different types of packets to different location thanks to the control plane abstraction. The DaMask-M contains two functions: *countermeasure selection* and *log generation*. When DaMask-M receives an alert, it tries to match the alert to a countermeasure. The default action is to drop the packet if there is no pre-set policy for the alert. We implement DaMask-M as a set of common APIs so that defenders can customize their own defense countermeasure for different DDoS attacks. The basic unit a defender can play with is flow. We define three basic operations, *forward*, *drop* and *modify* to form advanced defense logic. Compared with DDoS attack mitigation in traditional network, DaMask-M provides a powerful way to implement the countermeasure. After the countermeasure is selected, DaMask-M pushes the policy to the switch through network controller. After that, the attack packet, along with its auxiliary information (e.g. the time stamp and response actions), is recorded in the log database.

IV. DAMASK EVALUATION

We carried out a thorough performance evaluation of the DaMask architecture under various scenarios. We report the evaluation results in this section.

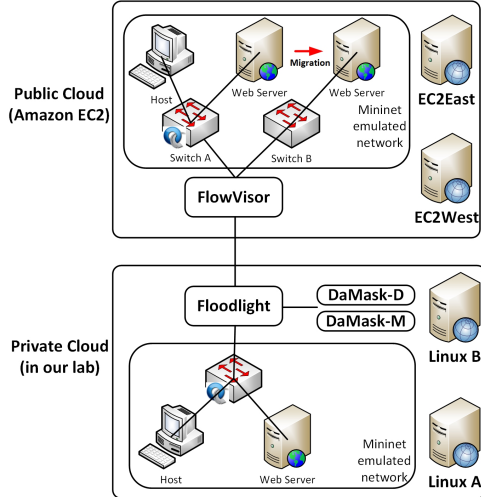


Fig. 3: The simulated hybrid cloud topology.

A. Evaluation Setting

To evaluate the performance of the DaMask, we have set up a hybrid cloud. We use Amazon Web Service EC2 as our public cloud while we simulate the private cloud in our lab. The overall evaluation environment is shown in Fig. 3. We utilize Mininet [9], which creates a realistic virtual network on a computer, to emulate the SDN setting.

1) *Private Cloud*: The private cloud consists of two Linux servers in our lab. Both of them are running Ubuntu 12.10 32-bit operating system. One laptop (denoted as Linux A), which equips with AMD E1-1200 \times 2 at 1.4 GHz CPU and 4 GB memory, emulates the private cloud. The other desktop (denoted as Linux B) equips with Intel i7-2600 CPU at 3.4 GHz and 12 GB memory runs the network controller and DaMask on it. Linux A and Linux B are connected through a Intel Express 460T 100MB switch. We choose Floodlight [10] as the network controller since Floodlight controller can be easily extended and enhanced through its module loading system.

We emulate a virtual network using Mininet in Linux A to extend the private cloud. The private cloud in Fig. 3 has one switch and two hosts. One of the hosts is an web server (Apache Http server 2.2.26). The Floodlight controller and the DaMask modules are deployed in Linux B. The DaMask modules communicate with the controller through Floodlight’s APIs.

2) *Public Cloud*: To measure the communication cost of DaMask, we use the Amazon Web Service (AWS) EC2 as our public cloud in our evaluation. We deployed two AWS EC2 instances as the company’s remote web servers. Both of them are Ubuntu T1-Micro instances. One of them (denoted as EC2West) is located at US West (Oregon) and the other (denoted as EC2East) is located at US East (N. Virginia).

¹The definition of the flow varies for each slice according to the specific requirements of each enterprise.

We use EC2West, which runs FlowVisor to handle network virtualization, to simulate the network administration of the public cloud. We emulate a virtual network in EC2East to extend the remote side of the company’s network, which is the public cloud part in Fig. 3. Similar with the private cloud extension, the extended public cloud also has one switch and two hosts, one of which is an Apache web server. The difference is that the switch is connected to the FlowVisor in EC2West instead of a network controller.

B. DaMask Overhead

TABLE I: Communication time

Task	Basic		DaMask w/o Test		DaMask w/ Test	
	West	East	West	East	West	East
Ping	196ms	12ms	425ms	51ms	462ms	85ms
Http	2.4s	1.7s	2.3s	1.6s	2.4s	1.6s

DaMask introduces communication overhead since now the traffic towards the servers in the public cloud needs to be examined by the DaMask-D module located at the enterprise’s local network. To evaluate the communication overhead, we carried out several experiments.

We first measured the network bandwidth of our evaluation environment. We measure the network bandwidth by running iperf 2 times a hour for a consecutive 24 hours. The average bandwidth between Linux A and EC2West is 27.4 MB/s, and the average bandwidth between Linux A and EC2East is 86.2 MB/s. The connection between Linux A and EC2East is better because our lab is located at east coast. We then tested the response time from the remote server with and without DaMask being deployed. We show our result in Table I. The results show that the communication overhead is only related to the round trip time between the server running the FlowVisor in the public cloud and the server running the network controller in the private cloud. This is because we fixed the size of the message to be sent to the network controller. Therefore, the communication overhead is a constant if the link status of network is stable.

Besides the communication overhead, the online detection algorithm will also introduce computation overhead. We embed Snort, one of the most popular open source detection system, along with Snort.AD, an anomaly-based detection plugin for Snort in the DaMask-D module. The online test process turns out to be quite efficient using our moderate off-the-shelf PC. The average inference time is 80 ms.

C. Adapting topology change

One advantage of DaMask is that DaMask is able to adapt the network topology change caused by virtual machine migrations. To simulate the migration process, we added an additional switch, i.e., switch B in Fig. 3. Suppose the web server is migrated from the switch A to the switch B, DaMask need taking control of the switch B while dropping control of the switch A since the switch A no longer belongs to the company’s slice.

Such re-configuration is accomplished by changing the flow space header of the company's slice in FlowVisor. Re-configuration in FlowVisor can be efficiently done through Command Line Interface (CLI) of FlowVisor. Since FlowVisor can reload the new slice configuration without interrupting the service, this process can be done in real-time.

After changing the flow space of the company's slice, we sent ICMP packets to the web server that is attached to the switch B. All the ICMP packets were received by the company's controller, which verifies the company indeed had the control of the switch B. We further tested if the company's control slice affected other users. We sent ICMP packets to the web server linked with the switch A and none of the packet was received by the company's controller, which means the FlowVisor did not forward any ICMP packet to the company's network controller.

V. RELATED WORK

Defending DDoS attack in traditional network has been studied for several decades. The surveys [11], [12] have included most of these work. Although our objective shares the similarity with them which is to defend DDoS attacks, our network environment which involves cloud computing and SDN is quite different from theirs. SDN technique has been used to address various network security. Jafarian et al. [13] proposed a random host mutation scheme using OpenFlow to achieve transparent moving target defense in SDN. Porras et al. proposed a security enforcement kernel for SDN in [14] to detect policy conflicts within the switches. Yao et al. utilized the SDN architecture to validate source addresses in [15]. The key difference between those work and ours is that they try to address the traditional network security threats using SDN to achieve better performance while we focus on the new challenges in the new network paradigm. Recently, Shin et al. [16] proposed an OpenFlow security application development framework, FRESCO, to enhance the secure application development in SDN. In contrast with FRESCO, our work focuses on DDoS attack challenges in cloud computing, which requires additional functionalities such as letting enterprises control the network slice other than those provided by existing solutions.

VI. CONCLUSION

Cloud computing is already here to stay and SDN is gaining increased popularity. With both of the technology emerging as the future enterprise IT solutions, it is worthwhile to look at the implications of the combination of the two, particularly on the enterprise network security. In this paper, we analyze the impact of cloud computing and SDN on DDoS attack defense. Based on our analysis, we identify the challenges and the benefits raised by these new technologies. We claim that with careful design, SDN could help with DDoS attack protection. To substantiate our finding, we proposed our solution of defending DDoS attack—DaMask architecture. Compared to the existing solutions, DaMask requires little effort from the cloud

provider which means few changes are required from the current cloud computing service architecture. The SDN-based network monitoring and control mechanism allow companies to control and configure their defense mechanisms in the cloud effectively without affecting other cloud users. We also carried out a simulation study using real network traces to evaluate the performance. The results show that our proposed DaMask is successful in dealing with the new challenges raised. The SDN-based network management can rapidly adapt to the network topological changes.

ACKNOWLEDGMENT

This work was supported in part by U.S. National Science Foundation under grant CNS-1217889.

REFERENCES

- [1] D. Geneiatakis, G. Portokalidis, and A. D. Keromytis, "A multilayer overlay network architecture for enhancing ip services availability against dos," in *7th International Conference, ICISS*, Kolkata, India, December 2011.
- [2] X. Liu, X. Yang, and Y. Lu, "To filter or to authorize: Network-layer dos defense against multimillion-node botnets," in *ACM SIGCOMM Computer Communication Review*. ACM, 2008.
- [3] P. Mittal, D. Kim, Y. C. Hu, and M. Caesar, "Mirage: Towards deployable ddos defense for web applications," August 2012.
- [4] W. G. Morein, A. Stavrou, D. L. Cook, A. D. Keromytis, V. Misra, and D. Rubenstein, "Using graphic turing tests to counter automated ddos attacks against web servers," in *10th ACM conference on CCS*. ACM, 2003.
- [5] H. Wang, L. Xu, and G. Gu, "Of-guard: A dos attack prevention extension in software-defined networks."
- [6] D. Kreutz, F. Ramos, and P. Verissimo, "Towards secure and dependable software-defined networks," in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*. ACM, 2013, pp. 55–60.
- [7] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live migration of virtual machines," in *Proceedings of the 2nd conference on Symposium NSDI*. USENIX, 2005.
- [8] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: enabling innovation in campus networks," *ACM SIGCOMM Computer Communication Review*, 2008.
- [9] B. Lantz, B. Heller, and N. McKeown, "A network in a laptop: rapid prototyping for software-defined networks," in *9th ACM SIGCOMM Workshop HotSDN*. ACM, 2010.
- [10] "Floodlight openflow controller," <http://floodlight.openflowhub.org>.
- [11] T. Peng, C. Leckie, and K. Ramamohanarao, "Survey of network-based defense mechanisms countering the dos and ddos problems," *ACM Comput. Surv.*, no. 1, April 2007.
- [12] S. Zargar, J. Joshi, and D. Tipper, "A survey of defense mechanisms against distributed denial of service (ddos) flooding attacks," *Communications Surveys Tutorials, IEEE*, 2013.
- [13] J. H. Jafarian, E. Al-Shaer, and Q. Duan, "Openflow random host mutation: transparent moving target defense using software defined networking," in *Proceedings of the 1st workshop on HotSDN, SIGCOMM*. ACM, 2012.
- [14] P. Porras, S. Shin, V. Yegneswaran, M. Fong, M. Tyson, and G. Gu, "A security enforcement kernel for openflow networks," in *Proceedings of the 1st workshop on HotSDN, SIGCOMM*. ACM, 2012.
- [15] G. Yao, J. Bi, and P. Xiao, "Source address validation solution with openflow/noxarchitecture," in *IEEE ICNP*, 2011.
- [16] S. Shin, P. Porras, V. Yegneswaran, M. Fong, G. Gu, and M. Tyson, "Fresco: Modular composable security services for software-defined networks," in *Proceedings of NDSS*, 2013.