This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

Exploring the firewall security consistency in cloud computing during live migration

PLEASE CITE THE PUBLISHED VERSION

<https://doi.org/10.1145/2967878.2967922>

PUBLISHER

ACM

VERSION

AM (Accepted Manuscript)

PUBLISHER STATEMENT

© ACM 2016. This is the author's version of the work. It is posted here for your personal use. Not for redistribution.

LICENCE

All Rights Reserved

REPOSITORY RECORD

Al-Amri, Shadha, and Lin Guan. 2019. "Exploring the Firewall Security Consistency in Cloud Computing During Live Migration". figshare. <https://hdl.handle.net/2134/23735>.

Exploring the Firewall Security Consistency in Cloud Computing during Live Migration

Shadha Mohamed ALAmri and Lin Guan

Department of Computer Science
Loughborough University
United Kingdom

S.AL-Amri@lboro.ac.uk, l.guan@lboro.ac.uk

ABSTRACT

Virtualization technology adds great opportunities and challenges to the cloud computing paradigm. Resource management can be efficiently enhanced by employing Live Virtual Machine Migration (LVMM) techniques. Based on the literature of LVMM implementation in the virtualization environment, middle-boxes such as firewalls do not work effectively after LVMM as it introduces dynamic changes in network status and traffic, which may lead to critical security vulnerabilities. One key security hole is that the security context of the firewall do not move with the Virtual Machine after LVMM is triggered. This leads to inconsistency in the firewall level of protection of the migrated Virtual Machine. There is a lack in the literature of practical studies that address this problem in cloud computing platform. This paper demonstrates a practical analysis using OpenStack testbed to study the firewalls limitations in protecting virtual machines after LVMM. Two network scenarios are used to evaluate this problem. The results show that the security context problem does not exist in the stateless firewall but can exist in the stateful firewall.

Categories and Subject Descriptors

C.2.0 [General]: Security and protection, firewall; C.2.1 [Network Architecture and Design]: Distributed networks

Keywords

cloud computing, live migration, firewall, OpenStack

1. INTRODUCTION

Cloud computing paradigm introduces an efficient utilisation of huge computing resources among multiple users with minimal expense and deployment effort compared to traditional computing facilities. However, since cloud computing has emerged from a business perspective, it has faced

the challenges of standardisation and compatibility[1]. Although cloud computing has incredible benefits, some governments and enterprises are hesitating to transfer their computing technology to the cloud as a consequence of security challenges. Security is, therefore, a significant factor in motivating the cloud computing adoption. Cloud services consist of three layers: Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). IaaS is the bottom layer where all the other layers, PaaS and SaaS are built on it. Technical security challenges emerge in different layers, but the focus of this paper is on IaaS. [2, 3, 4]

Since cloud computing infrastructure has the features of rapid elasticity and fast resource pooling, its services and processes are built on virtualization technology, which supports these features. One of the main core management processes in IaaS are Virtual Machine (VM) provisioning and VM migration, which are parts of virtualization technology [5]. There are two types of VM migration: live VM migration and non-live VM migration [6]. This paper focuses on live VM migration as it has significant advantages, such as enhancement of the performance of High Computing Performance (HCP)[7], minimising the downtime for the system with minimum human interaction[8], server consolidation in order to avoid sprawl, balancing workloads between physical machines to avoid large discrepancies among physical machines and mitigating hot spots when extra resources are needed [9].

Live Virtual Machine Migration (LVMM) is a technique that allows Virtual Machines to be moved at run-time between physical machines without interrupting their processes. However, LVMM utilisation has been restricted, as it entails security consequences. One major security vulnerability is that the security context of the VM does not move with it after migration. Therefore, the VM will not be protected by the firewall rules that were being set before migration took place. This reduces the efficiency of the firewalls and may allow unauthorised parties to access the VM.

This paper investigates this security problem by examining the behavior of the firewall before and after LVMM triggered in a cloud, as this problem has a lack of practical exploration in the literature in a cloud-based environment.

This paper is organised as follows: an overview of the work related to LVMM security holes is illustrated in Section 2. Section 3 examines the firewall consistency during live virtual machine migration. Section 4 contains the Discussion section. Finally, Section 5 presents the conclusions from this paper and future work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

2. RELATED WORK

There are two main categories of LVMM security vulnerabilities that have been identified in the literature. The first vulnerability is through a hole where the attackers can attack the VM running on the cloud system [10, 11, 12], since LVMM protocol transfer data in plaintext through network links. As a result, LVMM can be penetrated practically via man-in-the-middle attack [13, 14]. The second vulnerability is due to dynamic VM network state transfer, which affects transferring the security context with the VM after LVMM. The second category is the focus of this paper, as it faces a lack of practical investigation where the first category will be under examination for future work.

Security context consistency has been studied by [15] who presented a LVMM framework using NSE-H (Network Security Enabled Hypervisor) and built a prototype based on a stateful firewall in Xen hypervisor. As well, by [16] who implemented a framework using Java to migrate the security state along with the VM during LVMM in Linux with KVM hypervisor. However, their experiments did not run on IaaS platform.

Based on structural congruence and a reduction relation, [17] presented an algebraic framework named Cloud Calculus, which can help to design a topology of cloud computing with a firewall, where the security policies can be specified along with LVMM. Verification of this method was accomplished by translating the firewall configuration into CSP syntax and using Sugar SAT-solver [18]. However, this proposed work aims to verify the consistency of sets of firewall rules but does not aim to maintain a firewall configuration after LVMM takes place.

An analytical experiment in a real cloud, done by [19] in 2013, shows that there is a time interval where LVMM is not under firewall protection, but this study does not investigate the security context movement. An analysis of the virtual machine migration published by [20] which points out that the root security problem in LVMM is due to the change of IP address after migration leads to inconsistency of IP address which is configured in the firewall and IPS. However, the methodology used in setting up the experiment did not involve a cloud environment.

According to our best knowledge and investigation of the literature, the problem of firewall security policy consistency in VMs after the LVMM process is enabled in cloud computing has not been in a cloud environment. Therefore, this paper aims to investigate this problem in a cloud environment based on OpenStack cloud.

3. OPENSTACK IAAS CLOUD

In this section, an overview of OpenStack deployment has been illustrated and the used network topology has been presented. The main two components in this paper are the firewall and the Live Virtual Machine Migration. They have been defined and their configuration has been introduced within the scope of this paper.

3.1 OpenStack test-bed deployment

OpenStack is an open source cloud computing platform project that offers IaaS. OpenStack has been selected as the experiment platform in this paper because it has a very supportive community of both academic researchers and commercial bodies

In this paper, Ubuntu 14.04 is used as an operating system (OS) for all nodes as Ubuntu supports OpenStack cloud. Juno OpenStack release has been used to install the required packages for this cloud platform which consists of Controller node, Network node and Compute node. Figure1 shows the OpenStack services installed on each node.

Network node runs basic networking services which are DHCP Agent to provide DHCP service, L3 agent to provide routing service, L2 agent to provide virtual switching service and neutron-plugin-openvswitch-agent to provide openvswitch service.

Controller node is managing cloud infrastructure as it runs several services to manage the communication between different OpenStack cloud components via nova-api; where Keystone is managing authorization and authentication for services and users on this platform. Horizon is used to implement a dashboard that used to control most of the cloud operations via a web-based interface. Moreover, the controller is holding the network server configuration which installed via Neutron server service. Glance is used for managing operating system images that used during VM creation.

The main role for Compute nodes is to host virtual machines (VM) which are created on the cloud system using Nova-Compute. Compute nodes run as well L2 agent and neutron-plugin-openvswitch-agent in order to manage network connections between VMs[21]. The Controller is configured to have an extra role so it can also operate as a Compute node since there is a need to have two Compute nodes in order to implement LVMM.

The hardware specification for each physical server used in the experiments is 16GB RAM, 300GB hard disk, two network cards and Intel Core i5-4460 CPU@3.2GHz ×4.

Figure 1: Openstack services for each node

3.2 Network Topology

Networking in OpenStack uses either Nova-network or Neutron Network, which are types of SDN (Software Defined Network). Neutron is more flexible as it allows plug-in components such as OpenFlow and firewalls. Moreover, it is newer than nova-network; however, it is more complex [21]. As a first stage, a simple single flat network was configured, which consists of two networks: a Management Network used to manage the OpenStack components on the administration side, and a Data Network used to exchange data between VM (Virtual Machines). In terms of network de-

vices, one physical router-switch is used and one eight-port switch. Moreover, OVS (Open vSwitch) is another component of the SDN network.

The Management Network is used by the physical nodes for system control and management of traffic, whereas the Data Network is specifically for VM traffic only. This topology was configured using the Neutron Network service, and the networking component was installed on each node according to its functions as illustrated in Figure 2.

- Controller Node has a Neutron server configuration
- Compute Node has Layer2 agent, Open vSwitch agent
- Network Node has a DHCP agent, Lear3 agent, Meta-data agent and Open vSwitch agent.

In this experimental setting, the Controller Node acts as Controller and Compute2, therefore, it has the Compute Node network configuration as well.

Figure 2: Network connection between the Experiment servers

3.3 Firewall configuration

The Firewall task is to inspect the network packets and decide to allow them to pass or drop them based on pre-defined rules[22]. Based on how a decision is made for every packet, firewalls are categorised into stateless firewalls and stateful firewalls. If a firewall decides the fate of every packet solely by examining the packet itself, then the firewall is called a stateless firewall. If a firewall decides the fate of some packets not only by examining the packet itself, but also by examining the packets that the firewall has accepted previously, then the firewall is called a stateful firewall [23].

An iptables firewall is used for the testing. Virtual firewalls are preferable to traditional firewalls in cloud computing. Traditional firewall placement is not sufficient in cloud computing, as it will introduce traffic overhead to the network switches and hypervisors [24].

3.4 LVMM technique implementation

The ideal migration process is to copy the complete state of VM, including memory, disk and network connection [8]. There is a default LVMM algorithm in the most popular VMM such as Xen, VMWare and KVM. In order to migrate memory state; there are two mechanisms which are pre-copy and post-copy as compared by [25] In [26] illustrate the Sequential steps of the Xen default VM live migration algo-

rithm as displayed in Figure 3. This algorithm is based in transferring the run-time memory state of the VM.

Figure 3: Xen LVMM algorithm

Three methods can be used to implement LVMM in OpenStack: shared storage live migration that requires a shared storage system to be configured such as NFS, block live migration or volume-backed live migration that requires to install and configure an extra OpenStack service called cinder. In this experiment, block live migration is implemented as it meets the scope of this paper. It does not use a shared folder between the Compute nodes to store VMs and does not require an extra service to be installed LVMM is disabled by default in OpenStack, as it incurs security vulnerabilities. It is, therefore, required to configure the nova.conf files of the Compute nodes to enable live migration flag. It has been noticed that the CPU version on the nodes can affect LVMM process where the VM can not be migrated from a compute node that has a lower CPU version to another compute node that has a higher CPU version. Moreover, Libvirt has to be configured to listen for unsecured TCP connections. It is obvious from these configurations, that some security holes are opened in order to enable LVMM; however, these particular security aspects are under investigation.[21]

4. EXPERIMENT

In these experiments, two physical networks are created. The first network for management and used for control plane traffic. The second network for Data and used by VMs to communicate with each other. LVMM traffic takes place on Data network. In order to perform LVMM, two compute nodes are set-up. For testing purposes, three VMs have been launched: one in Compute1 and two in Compute2 as shown in Table 1. All VMs are in the same subnet of Data network.

Table 1: Scenario1: VMs locations

VM name	IP	location
vmtest1	19.168.1.33/24	Compute2
vmtest2	192.168.1.34/24	Compute1
vmtest3	192.168.1.35/24	Compute2

Problem definition. VM gets a dynamic IP address from a DHCP server, therefore, it is more likely that VM gets a different IP address when it shut down and/or DHCP leased line expired. In LVMM process, there is a period of time where the VM is not seen in the network link called down-time interval. Therefore, VM might get a different IP address after the migration process is completed.

Hypothesis. If LVMM is triggered, then VM IP address might change which affect the firewall rules consistency as firewall uses a static IP address to create its rules

Scenario no.1: Testing IP address Consistency. The aim of this scenario is to analysis if the VM's IP

address might change after live migration take place since VM got dynamic IPs from DHCP server.

DHCP leased line has been set to be one hour in order to allow it to expire during the LVMM process. The experiment in this scenario has been repeated once on a weekly basis for two months. The VM has been restarted many times due to LVMM down-time interval. Table 2 shows a code written on Ubuntu that used to repeat live migration process ten times per week.

Table 2: code to repeat LVMM

```
for (( i=1; i <= 5 ; i++ ))
do
nova live-migration --block-migrate 58aabbbea-e309-
4988-9b9e-fec99bb1a69a compute1
sleep 3m
nova show 58aabbbea-e309-4988-9b9e-fec99bb1a69a |
grep hypervisor
sleep 1m
nova live-migration --block-migrate 58aabbbea-e309-
4988-9b9e-fec99bb1a69a controller sleep 3m nova
show 58aabbbea-e309-4988-9b9e-fec99bb1a69a | grep
hypervisor
sleep 1m
done
```

Results. The results show that the VM IP did not change; even though the DHCP leased line expiration duration is due to. It has been found that most hypervisors make use of the ARP protocol in order to allow VM to keep their IP address after migration which agrees with the results obtained in [27, 28, 29]. While these results contradict with the methodology used in [20] where the authors assume that shutting down VM can mimic LVMM process and allow VMs to change their IP address which affects the security consistency of IPS and IDS systems.

Scenario no.2 : Testing Stateless Firewall Rules Consistency. This scenario has been conducted into two different network topologies. The first topology, all the VMs are on the same LAN. The second topology, the VMs separated into two VLANs. The aim of this setting is to study if the network topology introduces any effects on the stateless firewall consistency as VLAN introduce logical network isolation.

OpenStack platform uses a security group mechanism to set security rules and then neutron saves the rules in iptables. In this experiment, two basic rules are configured for verification purposes. The first rule restricts the incoming ICMP echo requests sent to a specific VM as illustrated in Table 3. The second rule restricts incoming TCP request through an ssh connection as illustrated in Table 4.

Table 3: firewall rule no.1

```
#iptables -A INPUT -p icmp --icmp-type echo-
request -s 192.168.1.33 -d 192.168.1.35 -m state --state
NEW,ESTABLISHED,RELATED -j DROP
```

The first part of this scenario is to trigger LVMM while all VMs in the same LAN. The second part of this scenario, two VLANs.

Results. The results of this scenario illustrate that the stateless firewall rules were valid before and after live migra-

Table 4: firewall rule no.2

```
#iptables -A INPUT -p --tcp -d 192.168.1.33 -s
192.168.1.34 -p tcp --dport 22 -j DROP
```

tion regardless if the VMs are all in one LAN or separated into different VLANs. **Finding.** The results of scenario no.1 and scenario no.2 show that the Hypothesis. is not valid. Live Virtual Machine Migration did not contribute in changing the VM IP address. As a consequence, stateless firewall consistency has not been affected.

5. DISCUSSION

Based on the experiments which involve configuration of two different security rules on two different network topologies, it has been found that the security context set on VM according to the stateless firewall rules are valid even after live virtual machine migration is triggered.

This experimental investigation leads to a critical analysis of various aspects in enhancing security of LVMM in cloud computing involving firewall as demonstrated below

- The security context consistency problem exists in the stateful firewall as has been explored by [30] as only packets matching a known connection state are allowed to go through the firewall. However, the proposed solution is based on a specific type of virtual firewalls which is vShield.
- A major network security risk in cloud computing is due to the limits of traditional firewall connections [31]. Moreover, traditional firewall settings are not sufficient for optimal fine-grained decisions, and application-level firewalls are not able to deal with dynamically opened server ports for encrypted connections [32]. In addition [30] explained the reasons for blind spots which are not covered by traditional security appliances and show that traditional firewalls cannot protect cloud traffic. Moreover, they point out that virtual firewalls can overcome this limitation.
- The most critical point is that security group technique is used by IaaS cloud to set the firewall security rules, but it lacks to many features that are provided by traditional firewalls [33] and make the firewall configuration more complex and trickier [34]. As well as it is a type of stateless firewall which does not support high dynamic environment such as cloud [30]. However, replacing stateless firewall with stateful firewall introduce the security context consistency problem. In fact, stateful firewall is still an open problem [35]

Therefore, the following aspects require further practical investigation in the cloud environment:

- In a context of the stateful firewall; the security context configured for VM is not moving with the corresponding VM after LVMM is triggered even though all VMs are located on the same LAN. That leads to nullifying the configured security policies.
- Live migration effects on QoS when security mechanisms are implemented to mitigate security holes opened during LVMM configuration.

6. CONCLUSIONS AND FUTURE WORK

In this paper, a firewall exploration has been studied to evaluate its security context during Live Virtual Machine migration. This security context was configured via a stateless firewall and the experiment was conducted on a cloud computing platform deployed using OpenStack. This experiment has shown that the authorization rules configured in a stateless firewall to protect access to VMs are consistent, even after Live Virtual Machine Migration (LVMM) is triggered. However, in cloud IaaS, the stateless firewall in the form of security group faces several limitations as it is complex to configure and it is considered as a type of traditional firewalls. Therefore, the proper type of firewall has a great impact on preserving security and maintaining the cloud resource management flexibility. This experimental evaluation in cloud testbed is a step in investigating the security challenges in LVMM in order to provide a practical and visible mitigation. Future work is planned to enhance the security of inter-communication in Infrastructure as a Service (IaaS) during Live Virtual Machine Migration (LVMM) when the stateful firewall is configured.

7. REFERENCES

- [1] Ian Foster, Yong Zhao, Ioan Raicu, and Shiyong Lu. Cloud computing and grid computing 360-degree compared. In *Grid Computing Environments Workshop, 2008. GCE'08*, pages 1–10. Ieee, 2008.
- [2] Chirag Modi, Dhiren Patel, Bhavesh Borisaniya, Avi Patel, and Muttukrishnan Rajarajan. A survey on security issues and solutions at different layers of cloud computing. *The Journal of Supercomputing*, 63(2):561–592, 2013.
- [3] M Azua Himmel and F Grossman. Security on distributed systems: Cloud security versus traditional it. *IBM Journal of Research and Development*, 58(1):3–1, 2014.
- [4] Luis M Vaquero, Luis Roderio-Merino, and Daniel Morán. Locking the sky: a survey on iaas cloud security. *Computing*, 91(1):93–118, 2011.
- [5] Rajkumar Buyya, James Broberg, and Andrzej M Goscinski. *Cloud computing: Principles and paradigms*, volume 87. John Wiley & Sons, 2010.
- [6] Raja Wasim Ahmad, Abdullah Gani, Siti Hafizah Ab Hamid, Muhammad Shiraz, Feng Xia, and Sajjad A Madani. Virtual machine migration in cloud data centers: a review, taxonomy, and open research issues. *The Journal of Supercomputing*, pages 1–43, 2015.
- [7] Viktor Mauch, Marcel Kunze, and Marius Hillenbrand. High performance cloud computing. *Future Generation Computer Systems*, 29(6):1408–1416, 2013.
- [8] Violeta Medina and Juan Manuel García. A survey of migration mechanisms of virtual machines. *ACM Computing Surveys (CSUR)*, 46(3):30, 2014.
- [9] Mayank Mishra, Anwesha Das, Purushottam Kulkarni, and Anirudha Sahoo. Dynamic resource management using virtual machine migrations. *Communications Magazine, IEEE*, 50(9):34–40, 2012.
- [10] L YamunaDevi, P Aruna, D Sudha Devi, and N Priya. Security in virtual machine live migration for kvm. In *Process Automation, Control and Computing (PACC), 2011 International Conference on*, pages 1–6. IEEE, 2011.
- [11] Wesam Dawoud, Ibrahim Takouna, and Christoph Meinel. Infrastructure as a service security: Challenges and solutions. In *Informatics and Systems (INFOS), 2010 The 7th International Conference on*, pages 1–8. IEEE, 2010.
- [12] Fengzhe Zhang, Yijian Huang, Huihong Wang, Haibo Chen, and Binyu Zang. Palm: security preserving vm live migration for systems with vmm-enforced protection. In *Trusted Infrastructure Technologies Conference, 2008. APTC'08. Third Asia-Pacific*, pages 9–18. IEEE, 2008.
- [13] Jon Oberheide, Evan Cooke, and Farnam Jahanian. Empirical exploitation of live virtual machine migration. In *Proc. of BlackHat DC convention*. Citeseer, 2008.
- [14] Xinyu Zhang, Yongli Zhao, Xin Su, Ruiying He, Weiwei Wang, and Jie Zhang. Load balancing algorithm based virtual machine dynamic migration scheme for datacenter application with optical networks. In *Communications and Networking in China (CHINACOM), 2012 7th International ICST Conference on*, pages 271–276. IEEE, 2012.
- [15] Chen Xianqin, Wan Han, Wang Sumei, and Long Xiang. Seamless virtual machine live migration on network security enhanced hypervisor. In *Broadband Network & Multimedia Technology, 2009. IC-BNMT'09. 2nd IEEE International Conference on*, pages 847–853. IEEE, 2009.
- [16] Zahra Tavakoli, Sebastian Meier, and Alexander Vensmer. A framework for security context migration in a firewall secured virtual machine environment. In *Information and Communication Technologies*, pages 41–51. Springer, 2012.
- [17] Yosr Jarraya, Arash Eghtesadi, Mourad Debbabi, Ying Zhang, and Makan Pourzandi. Cloud calculus: Security verification in elastic cloud computing platform. In *Collaboration Technologies and Systems (CTS), 2012 International Conference on*, pages 447–454. IEEE, 2012.
- [18] Yosr Jarraya, Arash Eghtesadi, Mourad Debbabi, Ying Zhang, and Makan Pourzandi. Formal verification of security preservation for migrating virtual machines in the cloud. In *Stabilization, Safety, and Security of Distributed Systems*, pages 111–125. Springer, 2012.
- [19] Mahwish Anwar. Virtual firewalling for migrating virtual machines in cloud computing. In *Information & Communication Technologies (ICICT), 2013 5th International Conference on*, pages 1–11. IEEE, 2013.
- [20] Beaulah Navamani, Chuan Yue, Xiaobo Zhou, and Edward Chow. An analysis of the virtual machine migration incurred security problems in the cloud. 2014.
- [21] Tom Fifield, Diane Fleming, Anne Gentle, Lorin Hochstein, Jonathan Proulx, Everett Toews, and Joe Topjian. *OpenStack Operations Guide*. " O'Reilly Media, Inc.", 2014.
- [22] Dirk Achenbach, Jörn Müller-Quade, and Jochen Rill. Universally composable firewall architectures using trusted hardware. In *International Conference on Cryptography and Information Security in the Balkans*, pages 57–74. Springer, 2014.
- [23] Mohamed G Gouda and Alex X Liu. A model of

- stateful firewalls and its properties. In *Dependable Systems and Networks, 2005. DSN 2005. Proceedings. International Conference on*, pages 128–137. IEEE, 2005.
- [24] Masoud Moshref, Minlan Yu, Abhishek Sharma, and Ramesh Govindan. vcrib: Virtualized rule management in the cloud. In *Proc. NSDI*, 2013.
- [25] Takahiro Hirofuchi, Hidemoto Nakada, Satoshi Itoh, and Satoshi Sekiguchi. Reactive consolidation of virtual machines enabled by postcopy live migration. In *Proceedings of the 5th international workshop on Virtualization technologies in distributed computing*, pages 11–18. ACM, 2011.
- [26] Sebastian Biedermann, Martin Zittel, and Stefan Katzenbeisser. Improving security of virtual machines during live migrations. In *Privacy, Security and Trust (PST), 2013 Eleventh Annual International Conference on*, pages 352–357. IEEE, 2013.
- [27] Wenjin Hu, Andrew Hicks, Long Zhang, Eli M Dow, Vinay Soni, Hao Jiang, Ronny Bull, and Jeanna N Matthews. A quantitative study of virtual machine live migration. In *Proceedings of the 2013 ACM Cloud and Autonomic Computing Conference*, page 11. ACM, 2013.
- [28] Qin Li, Jinpeng Huai, Jianxin Li, Tianyu Wo, and Minxiong Wen. Hypermip: Hypervisor controlled mobile ip for virtual machine live migration across networks. In *High Assurance Systems Engineering Symposium, 2008. HASE 2008. 11th IEEE*, pages 80–88. IEEE, 2008.
- [29] Petter Svärd, Benoit Hudzia, Steve Walsh, Johan Tordsson, and Erik Elmroth. Principles and performance characteristics of algorithms for live vm migration. *ACM SIGOPS Operating Systems Review*, 49(1):142–155, 2015.
- [30] Debashis Basak, Rohit Toshniwal, Serge Maskalik, and Allwyn Sequeira. Virtualizing networking and security in the cloud. *ACM SIGOPS Operating Systems Review*, 44(4):86–94, 2010.
- [31] Diogo AB Fernandes, Liliana FB Soares, João V Gomes, Mário M Freire, and Pedro RM Inácio. Security issues in cloud environments: a survey. *International Journal of Information Security*, 13(2):113–170, 2014.
- [32] Jan Wiebelitz, Michael Brenner, Christopher Kunz, and Matthew Smith. Early defense: enabling attribute-based authorization in grid firewalls. In *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*, pages 336–339. ACM, 2010.
- [33] Jordan Cropper, Johanna Ullrich, Peter Fruhwirt, and Edgar Weippl. The role and security of firewalls in iaas cloud computing. In *Availability, Reliability and Security (ARES), 2015 10th International Conference on*, pages 70–79. IEEE, 2015.
- [34] Cheng Jin, Anurag Srivastava, Yu Jin, and Zhi-Li Zhang. Secgras: Security group analysis as a cloud service. In *Network Protocols (ICNP), 2014 IEEE 22nd International Conference on*, pages 215–220. IEEE, 2014.
- [35] Yosr Jarraya, Arash Egtesadi, Sahba Sadri, Mourad Debbabi, and Makan Pourzandi. Verification of firewall reconfiguration for virtual machines migrations in the cloud. *Computer Networks*, 93:480–491, 2015.