Bays *et al. Journal of Internet Services and Applications* (2015) 6:1
DOI 10.1186/s13174-014-0015-z

Journal of Internet Services and Applications
a SpringerOpen Journal

RESEARCH Open Access

Virtual network security: threats, countermeasures, and challenges

Leonardo Richter Bays[1], Rodrigo Ruas Oliveira[1], Marinho Pilla Barcellos[1], Luciano Paschoal Gaspary[1*] and Edmundo Roberto Mauro Madeira[2]

Abstract

Network virtualization has become increasingly prominent in recent years. It enables the creation of network infrastructures that are specifically tailored to the needs of distinct network applications and supports the instantiation of favorable environments for the development and evaluation of new architectures and protocols. Despite the wide applicability of network virtualization, the shared use of routing devices and communication channels leads to a series of security-related concerns. It is necessary to provide protection to virtual network infrastructures in order to enable their use in real, large scale environments. In this paper, we present an overview of the state of the art concerning virtual network security. We discuss the main challenges related to this kind of environment, some of the major threats, as well as solutions proposed in the literature that aim to deal with different security aspects.

Keywords: Network virtualization; Security; Threats; Countermeasures

1 Introduction

Virtualization is a well established concept, with applications spanning several areas of computing. This technique enables the creation of multiple virtual platforms over a single physical infrastructure, allowing heterogeneous architectures to run on the same hardware. Additionally, it may be used to optimize the usage of physical resources, as an administrator is able to dynamically instantiate and remove virtual nodes in order to satisfy varying levels of demand.

In recent years, there has been a growing demand for adaptive network services with increasingly distinct requirements. Driven by such demands, and stimulated by the successful employment of virtualization for hosting custom-built servers, researchers have started to explore the use of this technique in network infrastructures. Network virtualization allows the creation of multiple independent virtual network instances on top of a single physical substrate [1]. This is made possible by instantiating one or more virtual routers on physical devices and establishing virtual links between these routers, forming topologies that are not limited by the structure of the physical network.

In addition to the ability to create different topological structures, virtual networks are also not bound by other characteristics of the physical network, such as its protocol stack. Thus, it is possible to instantiate virtual network infrastructures that are specifically tailored to the needs of different network applications [2]. These features also enable the creation of virtual testbeds that are similar to real infrastructures, a valuable asset for evaluating newly developed architectures and protocols without interfering with production traffic. [3] For these reasons, network virtualization has attracted the interest of a number of researchers worldwide, especially in the context of Future Internet research. Network virtualization has been embraced by the Industry as well. Major Industry players – such as Cisco and Juniper – nowadays offer devices that support virtualization, and this new functionality allowed infrastructure providers to offer new services.

In contrast to the benefits brought by network virtualization, the shared use of routing devices and communication channels introduces a series of security-related concerns. Without adequate protection, users from a virtual network might be able to access or even interfere with traffic that belongs to other virtual networks, violating

*Correspondence: paschoal@inf.ufrgs.br
[1] Institute of Informatics, Federal University of Rio Grande do Sul, Porto Alegre, Brazil
Full list of author information is available at the end of the article

security properties such as confidentiality and integrity [4,5]. Additionally, the infrastructure could be a target for denial of service attacks, causing availability issues for virtual networks instantiated on top of it [6,7]. Therefore, it is of great importance that network virtualization architectures offer protection against these and other types of threats that might compromise security.

Recently, attention has been drawn to security concerns in network infrastructures due to the discovery of pervasive electronic surveillance around the globe. Although all kinds of networks are potentially affected, the shared use of physical resources in virtual network environments exacerbates these concerns. As such, these recent circumstances highlight the need for a comprehensive analysis of current developments in the area of virtual network security.

In this paper, we characterize the current state of the art regarding security in network virtualization. We identify the main threats to network virtualization environments, as well as efforts aiming to secure such environments. For this study, an extensive literature search has been conducted. Major publications from the literature have been studied and grouped according to well known classifications in the area of network security, as well as subcategories proposed by the authors of this paper. This organization allows the analysis and discussion of multiple aspects of virtual network security.

The remainder of this paper is organized as follows. Section 2 presents a brief background on the area of network virtualization as well as a review of related literature. Section 3 introduces the taxonomy used to classify the selected publications. Section 4 exposes the security vulnerabilities and threats found in the literature, while Section 5 presents the security countermeasures provided by solutions found in previous proposals. In Section 6, we discuss the results of this study, and in Section 7 we summarize the main current research challenges in the area of virtual network security. Last, in Section 8 we present our conclusions.

2 Background and literature review

In this section, we first provide a brief background on the area of network virtualization, highlighting its most relevant concepts. Next, we present a review of literature closely related to virtual network security.

2.1 Background

Network virtualization consists of sharing resources from physical network devices (routers, switches, etc.) among different virtual networks. It allows the coexistence of multiple, possibly heterogeneous networks, on top of a single physical infrastructure. The basic elements of a network virtualization environment are shown in Figure 1. At the physical network level, a number of autonomous systems are represented by interconnected network substrates (e.g., substrates A, B, and C). Physical network devices are represented by nodes supporting virtualization technologies. Virtual network topologies (e.g., virtual networks 1 and 2), in turn, are mapped to a subset of nodes from one or more substrates. These topologies are composed of virtual routers, which use a portion of the resources available in physical ones, and virtual links, which are mapped to physical paths composed of one or more physical links and their respective intermediate routers.

From the point of view of a virtual network, virtual routers and links are seen as dedicated physical devices. However, in practice, they share physical resources with routers and links from other virtual networks. For this reason, the virtualization technology used to create this environment must provide an adequate level of isolation in order to enable the use of network virtualization in real, large scale environments.

Over the years, different methods for instantiating virtual networks have been used. Typical approaches include VLANs (Virtual Local Area Networks) and VPNs (Virtual Private Networks). Recently, Virtual Machine Monitors and programmable networks have been employed to create virtual routers and links over physical devices and communication channels. These approaches are briefly revisited next.

2.1.1 Protocol-based approaches

Protocol-based approaches consist of implementing a network protocol that enables the distinction of virtual networks through techniques such as tagging or tunneling. The only requirement of this kind of approach is that physical devices (or a subset of them) support the selected protocol.

One example of protocol-based network virtualization are VLANs. VLANs consist of logical partitions of a single underlying network. Devices in a VLAN communicate with each other as if they were on the same Local Area Network, regardless of physical location or connectivity. All frames sent through a network are tagged with their corresponding VLAN ID, processed by VLAN-enabled routers and forwarded as necessary [8]. Since isolation is typically based only on packet tagging, this approach is susceptible to eavesdropping attacks.

Another commonly used approach is the creation of Virtual Private Networks. VPNs are typically used to provide a secure communication channel between geographically distributed nodes. Cryptographic tunneling protocols enable data confidentiality and user authentication, providing a higher level of security in comparison with VLANs. VPNs can be provided in the physical, data link, or network layers according to the protocols being employed [9].

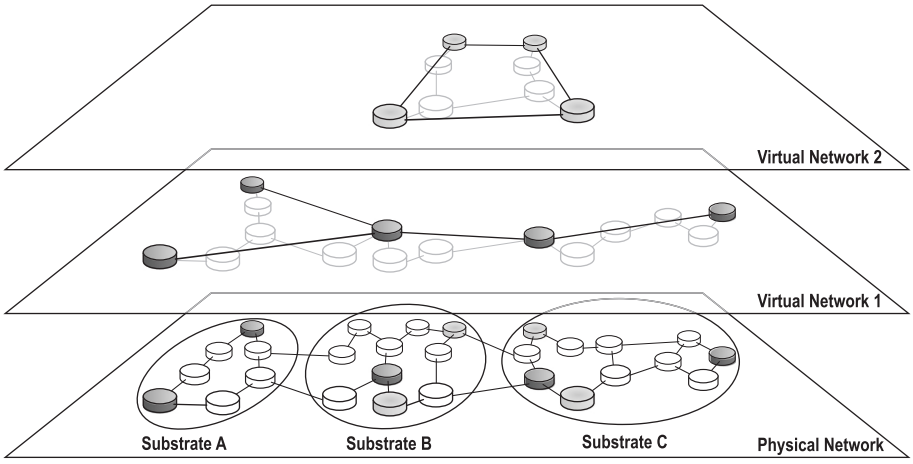

Figure 1 Network virtualization model, denoting a scenario with multiple physical substrates (Substrate A, B, and C) and virtual networks (Virtual Network 1 and 2).

2.1.2 Machine virtualization-based approaches

Machine virtualization-based approaches consist of creating virtual networks by means of groups of interconnected virtual machines. Virtual Machine Monitors are used to instantiate virtual routers, and virtual links are established between them, regardless of physical network topology. Table 1 shows different machine virtualization-based techniques that can be used to create virtual networks, as well as a brief explanation and an example of each.

This alternative is remarkably flexible and relatively cheap, as it allows the use of customized software and does not require the use of specific hardware[1]. However, it is more demanding in terms of resource usage in comparison to previously described protocol-based approaches. Additionally, it may introduce security concerns associated with server virtualization, some of which are mentioned in Sections 4 and 5. A general study on the security issues that arise from the use of machine virtualization was performed by van Cleeff *et al.* [10].

2.1.3 Programmable networks

Programmable routers have been used to enable the creation of virtual networks. Although this is not a new concept, research in this area has been recently stimulated by the inception of Software-Defined Networking (SDN). This paradigm consists of decoupling the data plane and the control plane in network devices. More specifically, devices such as routers and links retain only the data plane, and a separated control plane manages such devices based on an overview of the entire network.

OpenFlow [11], one of the most promising techniques for implementing this paradigm, defines a protocol that allows a centralized controller to act as the control plane, managing the behavior of network devices in a dynamic manner. The controller communicates with network devices through a secure connection, creating and managing flow rules. Flow rules instruct network devices on how to properly process and route network traffics with distinct characteristics. Through the establishment of specific flow rules, it is possible to logically partition physical networks and achieve data plane isolation. This isolation enables the creation of virtual networks on top of an SDN environment. OpenFlow gave rise to the Open Networking Foundation, an organization ran by major companies within the area of computer networks that aims to disseminate this type of technology.

Table 1 Virtualization techniques

Technique	Description	Examples
Full virtualization	The Virtual Machine Monitor emulates a complete machine, based on the underlying hardware architecture. The guest Operating System runs without any modification.	VMware Workstation, VirtualBox
Paravirtualization	The Virtual Machine monitor emulates a machine which is similar to the underlying hardware, with the addition of a hypervisor. The hypervisor allows the guest Operating System to run complex tasks directly on non-virtualized hardware. The guest OS must be modified in order to take advantage of this feature.	VMware ESX, Xen
Container-based virtualization	Instead of running a full Virtual Machine, this technique provides Operating System-level containers, based on separate userspaces. In each container, the hardware, as well as the Operating System and its kernel, are identical to the underlying ones.	OpenVZ, Linux VServer

2.2 Literature review

To the best of our knowledge, there have been no previous attempts at characterizing the state of the art regarding security in network virtualization. However, there have been a number of similar studies in other, closely related fields of research. We now proceed to a review of some of the main such studies.

Chowdhury *et al.* [1] provide a general survey in the area of network virtualization. The authors analyze the main projects in this area (both past projects and, at the time of publication, current ones) and discuss a number of key directions for future research. The authors touch upon the issues of security and privacy both while reviewing projects and discussing open challenges; however, as this is not the main focus of this survey, there is no in-depth analysis of security issues found in the literature.

Bari *et al.* [12] present a survey that focuses on data center network virtualization. Similarly to the aforementioned study, the authors survey a number of key projects and discuss potential directions for future work. When analyzing such projects, the authors provide insights on the fault-tolerance capabilities of each one, in addition to a brief discussion on security issues as one of the potential opportunities for future research.

In addition to the general studies on network virtualization presented so far, a number of surveys on cloud computing security have also been carried out. Cloud computing environments tend to make use of both machine and network virtualization, making this a highly relevant related topic for our study. However, while there is some overlap between cloud computing security and virtual network security, we emphasize that cloud computing represents a very specific use case of network virtualization and, therefore, poses a significantly distinct set of security challenges. Zhou *et al.* [13] provide an investigation on security and privacy issues of cloud computing system providers. Additionally, the authors highlight a number of government acts that originally intended to uphold privacy rights but fail to do so in light of advances in technology. Hashizume *et al.* [14], in turn, focus on security vulnerabilities, threats, and countermeasures found in the literature and the relationships among them.

Last, Scott-Hayward *et al.* [15] conducted a study on SDN security. As explained in Section 2.1.3, this is one of the technologies on top of which network virtualization environments can be instantiated. The authors first analyze security issues associated with the SDN paradigm and, afterwards, investigate approaches aiming at enhancing SDN security. Last, the authors discuss security challenges associated with the SDN model.

3 Taxonomy

The first step towards a comprehensive analysis of the literature was the selection of a number of publications from quality conferences and journals. To this end, we performed extensive searches in the ACM and IEEE digital libraries using a number of keywords related to network virtualization and security. We then ranked the literature found through this process according to the average ratio of citations per publication of the conferences or journals in which these papers were published. All publications from top tier conferences or journals with a consistent number of citations per publication were considered relevant and, therefore, selected. The remaining papers were analyzed and generally discarded.

Following the aforementioned process, a taxonomy was created in order to aid the organization and discussion of the selected publications. For this purpose, two well known classifications in the area of network security were chosen. Papers are organized according to the *security threats* they aim to mitigate, and afterwards, according to the *security countermeasures* they provide. As different authors have different definitions for each of these concepts, these classifications are briefly explained in the following subsections. The direct connection between them and the area of virtual network security is explained in sections 4 and 5, respectively.

In addition to these broad classifications, subcategories were created in order better organize this body of work. Figure 2 presents the full hierarchical organization that will be used in sections 4 and 5. Dark gray boxes represent broad categories used in the literature [16,17], while white boxes denote subdivisions proposed and created by the authors of this paper.

3.1 Security vulnerabilities and threats

There are a number of potential malicious actions, or threats, that may violate security constraints of computational systems. Shirey [16] describes and divides the consequences of these threats into four categories, namely *disclosure*, *deception*, *disruption*, and *usurpation*.

Unauthorized *disclosure* is defined as gaining unauthorized access to protected information. Sensitive data may be erroneously exposed to unauthorized entities, or acquired by an attacker that circumvents the system's security provisions.

Deception is characterized by intentionally attempting to mislead other entities. For example, a malicious entity may send false or incorrect information to others, leading them to believe that this information is correct. Fake identities may be used in order to incriminate others or gain illegitimate access.

Disruption means causing failure or degradation of systems, negatively affecting the services they provide.

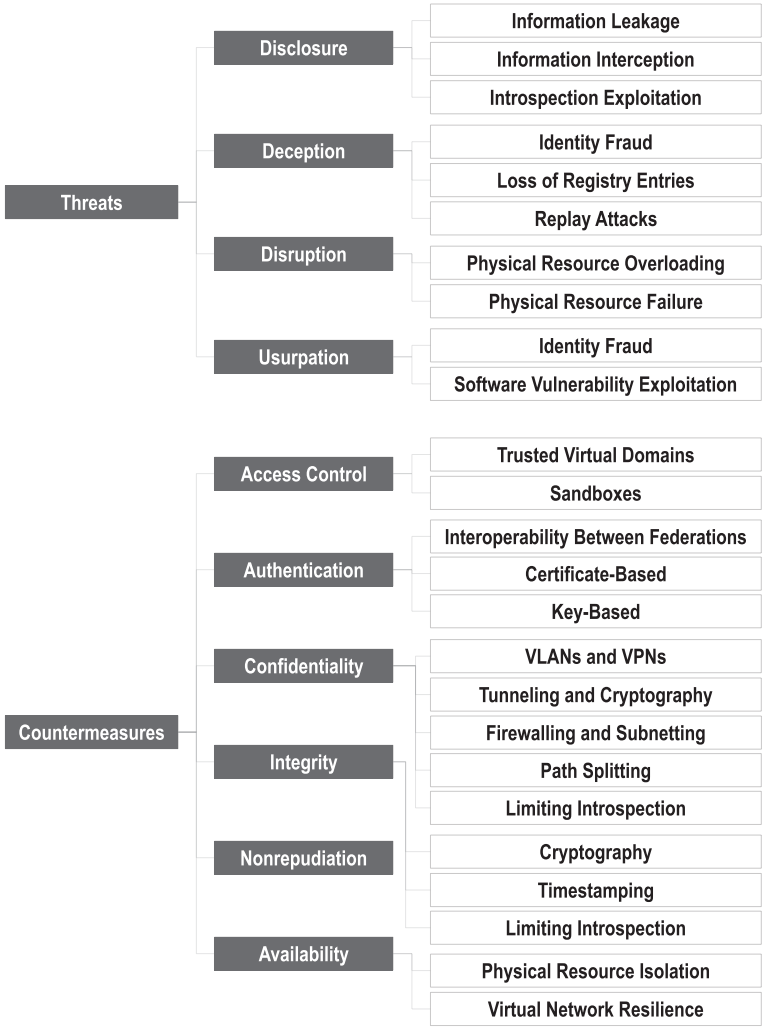

Figure 2 **Taxonomy used to classify publications in the area of virtual network security.**

This may be done by directly incapacitating a system component or the channel through which information is delivered, or by inducing the system to deliver corrupted information.

Last, through *usurpation*, an attacker may gain unauthorized control over a system. This unauthorized control may allow the attacker to illegitimately access protected data or services, or tamper with the system itself in order to cause incorrect or malicious behavior.

These threat categories, as well as the previously mentioned subcategories we have created, also cover vulnerabilities and attacks. For ease of comprehension, vulnerabilities and threats are discussed collectively in Section 4. Table 2 presents the relationships between vulnerabilities and threats in network virtualization environments. This table is organized according to the previously described taxonomy and lists all vulnerabilities found in the literature and the threats associated with each one. Additionally, the terms threat and attack are used interchangeably throughout the paper, as a threat may be understood as a potential attack (while an attack is the proper action that takes advantage of a vulnerability to violate a security policy).

3.2 Security countermeasures

Due to the existence of the previously described threats, computational systems must provide a series of countermeasures in order to maintain a desirable level of security. Stallings [17] categorizes these essential countermeasures into six subdivisions (referred to by Stallings as "security services"), namely *access control*, *authentication*, *data confidentiality*, *data integrity*, *nonrepudiation*, and *availability*.

Access control allows a system to administer which entities will be able to access its functions, and what permissions each of these entities will have. In order to grant

Table 2 Relationships between vulnerabilities and threats in network virtualization environments

Threat categories		Vulnerabilities	Threats
Disclosure	Information Leakage	Lack of ARP table protection	ARP table poisoning
		Placement of firewall rules inside virtual nodes	Subversion of firewall rules
	Information Interception	Lack of ARP table protection	ARP table poisoning
		Transmission of data in predictable patterns	Traffic Analysis attacks
		Uncontrolled handling of multiple, sequential virtual network requests from a single entity	Inference and disclosure of sensitive topological information
		Unprotected exchange of routing information among virtual routers	Disclosure of sensitive routing information
	Introspection Exploitation	Uncontrolled Introspection	Data theft
Deception	Identity Fraud	Improper handling of identities: - within individual networks;	Injection of malicious messages with forged sources
		- among federated networks;	Privilege escalation
		- during migration procedures.	Abuse of node removal and re-addition in order to obtain new (clean) identities
	Loss of registry entries	Uncontrolled rollback operations	Loss of registry entries
	Replay attacks	Lack of unique message identifiers	Replay attacks
Disruption	Physical Resource Overloading	Uncontrolled resource allocation	Performance degradation
			Abusive resource consumption
		Uncontrolled handling of virtual network requests	Exhaustion of resources in specific parts of the infrastructure
		Lack of proactive or reactive recovery strategies	Denial of Service attacks
	Physical Resource Failure	Lack of proactive or reactive recovery strategies	Failure of virtual routers/networks
		Uncontrolled resource reallocation after failures	Overloading of remaining virtual routers after failures
Usurpation	Identity Fraud	Improper handling of identities and associated privileges	Privilege escalation
	Software Vulnerability Exploitation	Privilege escalation in Virtual Machine Monitors	Unauthorized control of physical routers

individual access rights and permissions, entities must be properly authenticated in the system.

The purpose of *authentication* is to ensure that entities communicating with each other are, in fact, the entities they claim to be. The receiver of a message must be able to correctly identify its sender, and an entity must not be able to impersonate another.

Providing adequate *data confidentiality* means ensuring that third parties do not have access to confidential information being transmitted between two entities. Additionally, the system should inhibit attackers from deriving information by analyzing traffic flow characteristics.

Data integrity has the purpose of assuring that data stored by entities or transmitted through a network are not corrupted, adulterated or destroyed. Attacks such as duplication, modification, reordering, and replay of messages must be prevented. Furthermore, mechanisms for recovering from data corruption may also be provided.

In communications between peers, *nonrepudiation* provides a way to settle disputes when an entity denies having performed a certain action. The goal of this service is to prevent entities from falsely denying participation in any (possibly malicious) network-related activity.

The last security countermeasure is *availability*. System resources must be available upon request by an authorized entity, and the system must also conform to its performance specifications. In order to maintain availability, countermeasures against attacks such as *denial of service* must be provided.

4 Security vulnerabilities and threats

In this section, we present a comprehensive list of vulnerabilities and threats found in network virtualization environments. The interested reader should refer to Table 2 for a systematic review of such vulnerabilities and threats.

While some of the threats listed in this section are a result of accidental actions, we emphasize that all

threats – intentional or accidental – have an effect on security. As an example of an accidental attack, it is common for virtual routers to attempt to use all available resources (as virtualization tends to be transparent and virtual routers are typically not aware that they are not running on dedicated physical hardware). If the network virtualization environment does not adequately limit the resource usage of each virtual router, even this unintentional abuse may cause disruption on other networks hosted on the same substrate or cause the degradation or failure of critical services provided by the virtualization environment.

4.1 Disclosure

In an environment where physical resources are shared between a number of virtual networks, there is a series of behaviors that may result in undesired disclosure of information. Threats related to disclosure of private or sensitive information are explained next.

4.1.1 Information leakage

Cavalcanti *et al.* [18] mention the possibility of messages being leaked from one virtual network to another. In this type of attack, an entity may disclose private or sensitive information to members of other virtual networks, who should not have access to such information. The authors state that this may be achieved through ARP table poisoning. For example, a malicious user may spoof the IP address of a node that is able to send messages to the virtual network with which it intends to communicate. Wolinsky *et al.* [19] describe a similar attack, in which virtual nodes send messages to outside the boundaries of a network virtualization environment. This would make it possible for messages to reach physical nodes that not only do not belong to any virtual network, but are hosted outside of the virtualized network infrastructure. According to the authors, if data isolation is achieved by means of firewall rules, malicious users may be able to subvert such rules by escalating privileges and gaining root access on a virtual node.

4.1.2 Information interception

Attackers in a virtual network environment may capture messages being exchanged between two entities in order to access their content. This type of attack, often referred to as "eavesdropping" or "sniffing", may lead to theft of confidential information [4,5,20]. Wu *et al.* [20], specifically, mention ARP table poisoning as a means of achieving this. In contrast to the ARP poisoning attack described by Cavalcanti *et al.* [18] (explained in Section 4.1.1), in this case the attack would be used in order to mislead physical routers into forwarding packets meant to one entity to another one, allowing a malicious entity to sniff such packets. This is a common threat in any networking environment, but the use of shared physical resources by multiple virtual networks further exacerbates this problem. According to these and other authors, such as Cui *et al.* [21], networking solutions provided by virtual machine monitors may not properly isolate data belonging to different virtual networks. This means that members of one virtual network may be able to access data being transferred by other virtual networks sharing the same substrate.

Even if data inside network packets is protected (e.g. through the use of cryptography), entities may be able to derive sensitive information by analyzing them. In traffic analysis attacks, described by Huang *et al.* [22], entities acquire such information by analyzing characteristics of traffic flows between communicating entities in virtual networks. These characteristics include which entities communicate with which other entities, frequency of communication, and packet sizes, among others. For example, an entity that is involved in frequent, short communications with a high number of other entities may be a central point of control in the network. Knowing this, a malicious user could launch an attack directed at that entity, aiming to cause a considerable amount of disruption with limited effort. As previously mentioned, this attack is effective even if traffic is encrypted, making any type of virtual networking environment a potential target.

In addition to the previously detailed forms of information interception, which may also affect traditional network environments, other forms are specific to network virtualization. One such form is the use of multiple virtual network requests to disclose the topology of the physical infrastructure, explored by Pignolet *et al.* [23]. This constitutes a security threat, as infrastructure providers typically do not wish to disclose this information. The authors demonstrate that by sequentially requesting a number of virtual networks with varying topological characteristics and analyzing the response given by the infrastructure provider (*i.e.*, whether the request can be embedded or not), they are able to gradually obtain information about the physical topology. Moreover, the authors determine the number of requests needed to fully disclose the physical topology on networks with different topological structures (tree, cactus, and arbitrary graphs). Conversely, Fukushima *et al.* [24] state that the entity controlling a physical network may obtain confidential routing information from virtual networks hosted on top of it. As current routing algorithms require routing information to be sent and received through virtual routers, sensitive information may be disclosed to the underlying network.

4.1.3 Introspection exploitation

Introspection is a feature present in virtual machine monitors that allows system administrators to verify the current state of virtual machines in real time. It enables external

observers to inspect data stored in different parts of the virtual machine (including processor registers, disk, and memory) without interfering with it. While this feature has valuable, legitimate uses (*e.g.*, enabling administrators to verify that a virtual machine is operating correctly), it may be misused or exploited by attackers in order to access (and potentially disclose) sensitive data inside virtual machines [10]. This problem is aggravated by the fact that virtual nodes may be moved or copied between multiple virtual machine monitors, as sensitive data may be compromised through the exploitation of this feature on any virtual machine monitor permanently or temporarily hosting such virtual nodes.

4.2 Deception

We have identified three subcategories of threats that may lead to deception in virtual network environments. These subdivisions – namely identity fraud, loss of registry entries and replay attacks – are explained next.

4.2.1 Identity fraud

In addition to dealing with unauthorized disclosure, Cabuk *et al.* [5] and Wu *et al.* [20] also describe threats related to deception in virtual network environments. Specifically, virtual entities may inject malicious messages into a virtual network, and deceive others into believing that such messages came from another entity.

Certain characteristics of virtualized network environments increase the difficulty of handling identity fraud. The aggregation of different virtual networks into one compound network, known as federation, is indicated by Chowdhury *et al.* [25] as one of such characteristics. Federation raises issues such as the presence of separate roles and possible incompatibility between security provisions or policies from aggregated networks. Another complicating factor mentioned by the authors is the dynamic addition and removal of entities. An attacker may force a malicious node to be removed and re-added in order to obtain a new identity.

Other characteristics that complicate the handling of identity fraud involve operations such as migration and duplication of virtual nodes, as mentioned by van Cleeff *et al.* [10]. The study presented by the authors refers to virtualization environments in general. Therefore, in the context of this study, a virtual node may refer to either a virtual router or a virtual workstation. If a virtual node is migrated from one physical point to another, the identity of the machine that contains this virtual node may change. Moreover, virtual nodes may be copied to one or more physical points in order to provide redundancy, which may lead to multiple entities sharing a single identity. Both of these issues may cause inconsistencies in the process of properly identifying the origin of network messages, which may be exploited in identity fraud attacks.

4.2.2 Loss of registry entries

Van Cleeff *et al.* [10] also mention issues related to logging of operations in virtualization environments. If information regarding which entity was responsible for each operation in the network is stored in logs inside virtual machines, entries may be lost during rollback procedures. Likewise, logs of malicious activities performed by attackers may also be lost.

4.2.3 Replay attacks

Fernandes and Duarte [26] mention replay attacks as another form of deception in virtual networks. In this type of attack, a malicious entity captures legitimate packets being transfered through the network and retransmits them, leading other entities to believe that a message was sent multiple times. The authors explain that virtual routers may launch attacks in which they repeat old control messages with the intention of corrupting the data plane of the attacked domain.

4.3 Disruption

In a network virtualization environment, proper management of resources is crucial to avoid disruption. The main sources of disruption in such environments are related to the abuse of physical resources (either intentional or unintentional) and the failure of physical devices.

4.3.1 Physical resource overloading

Physical resource overloading may lead to failure of virtual nodes, or cause the network performance to degrade below its minimum requirements. This degradation may cause congestion and packet loss in virtual networks, as stated by Zhang *et al.* [27]. In addition to causing disruption in already established networks, overloading may also hinder the deployment of new ones.

Resource requirements themselves can be a point of conflict in virtual network environments. As explained by Marquezan *et al.* [28], multiple virtual networks may require an excessive amount of resources in the same area of the substrate network. While such prohibitive demands may be unintentional, they may also be due to a coordinated attack. This may not only happen during deployment operations, but also during the lifetime of virtual networks.

It is also possible for one virtual network to disrupt another by using more than its fair share of resources. This concern is explored by a number of authors in their respective publications [26,29-31]. Isolation and fair distribution of physical resources among virtual networks are essential to maintain the network virtualization environment operating properly. This includes assuring that the minimum requirements of each network will be fulfilled, as well as prohibiting networks from consuming more resources than they are allowed to.

Overloading may also be caused by attacks aimed at the physical network infrastructure. Attacks may originate from within a virtual network hosted in the same environment, or from outside sources. The most common threats are Denial of Service (DoS) attacks, as presented by Yu *et al.* [6] and Oliveira *et al.* [7]. A single physical router or link compromised by a DoS attack may cause disruption on several virtual networks currently using its resources.

4.3.2 Physical resource failure

As previously stated, the failure of physical devices is one of the sources of disruption in virtual infrastructures [32-34]. Possible causes range from the failure of single devices (a physical router, for example, may become inoperative if one of its components malfunctions) to natural disasters that damage several routers or links in one or more locations [35]. Additionally, further complications may arise as the remainder of the network may be overloaded during attempts to relocate lost virtual resources. In addition to being valuable from the point of view of fault tolerance, countermeasures for mitigating the effect of failures may also be applied in the event of attacks such as DoS, as in both cases there is a need for redirecting network resources away from compromised routers or links.

4.4 Usurpation

In virtual network environments, usurpation attacks may allow an attacker to gain access to privileged information on virtual routers, or to sensitive data stored in them. Such attacks may be a consequence of identity fraud or exploited vulnerabilities, which are explained next.

4.4.1 Identity fraud

As previously mentioned in Section 4.2, identity fraud attacks can be used to impersonate other entities within a virtual network. By impersonating entities with high levels of privilege in the network, attackers may be able to perform usurpation attacks. As an example, the injection of messages with fake sources mentioned by Cabuk *et al.* [5] is used for this purpose. By sending a message that appears to have been originated from a privileged entity, attackers may perform actions restricted to such entities, including elevating their own privilege level.

4.4.2 Software vulnerability exploitation

Roschke *et al.* [36] mention that virtual machine monitors are susceptible to the exploit of vulnerabilities in their implementation. According to the authors, by gaining control over a virtual machine monitor, attackers can break out of the virtual machine, obtaining access to the hardware layer. In an environment that uses full virtualization or paravirtualization to instantiate virtual routers, exploiting such vulnerabilities may enable an attacker to have full control over physical routers. By gaining access to physical devices, attackers could easily compromise any virtual networks provided by the infrastructure. As examples of such threats in practice, the Common Vulnerabilities and Exposures system lists a number of vulnerabilities in different versions of VMware products that allow guest Operating System users to potentially execute arbitrary code on the host Operating System [37-40].

5 Security countermeasures

In this section, we explore solutions published in the literature that aim to provide security and protect the environment from the aforementioned security threats.

5.1 Access control

Access control makes use of authentication and authorization mechanisms in order to verify the identity of network entities and enforce distinct privilege levels for each. This countermeasure is approached in two different manners in the literature, namely Trusted Virtual Domains and sandboxes. While these approaches are closely related to the notion of controlled execution domains, note that access control is performed in order to ensure that entities are granted the appropriate privilege levels.

5.1.1 Trusted virtual domains

Cabuk *et al.* [5] devised a framework to provide secure networking between groups of virtual machines. Their security goals include providing isolation, confidentiality, integrity, and information flow control in these networks. The framework provides the aforementioned security countermeasures through the use of Trusted Virtual Domains (TVDs). Each TVD represents an isolated domain, composed of "virtualization elements" and communication channels between such elements. In Cabuk's proposal, the virtualization elements are virtual workstations. However, the concept of TVDs may be applied to any device supporting virtualization.

Figure 3 depicts a virtual network infrastructure with three TVDs (A, B, and C). Gray routers represent gateways between these domains. While the gateway between TVDs B and C is simultaneously within both domains, the gateways between A and B are isolated – making use of an auxiliary TVD (AB) in order to communicate.

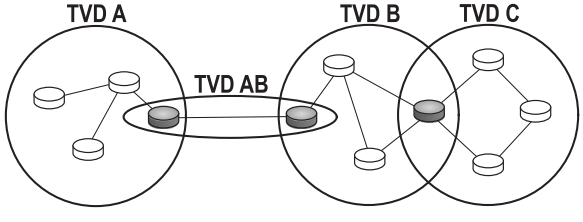

Figure 3 Example of a virtual infrastructure with three Trusted Virtual Domains, as described by Cabuk *et al.* [5].

Access control is performed when virtual machines join a TVD, ensuring that only machines that satisfy a given set of conditions are able to join. This admission control may be applied continuously in case prerequisites to join a TVD are changed. Additionally, TVDs leverage access policies to prevent unauthorized access.

5.1.2 Sandboxes

Wolinsky *et al.* [19] use virtual machine sandboxes in order to provide security in large scale collaborative environments. Although this work focuses on networked virtual machines hosting virtual workstations, this concept can be extended to virtual networks. Sandboxes are used to limit virtual machine access to physical resources, preventing malicious virtual machines from accessing data within other virtual machines. Moreover, each virtual machine supports IPSec, enabling the creation of secure communication channels, and X.509, providing virtual machine authentication. The authentication process is detailed in Section 5.2.

5.2 Authentication

Authentication aims to ensure that entities in a network environment are who they claim to be. In virtual network environments, providing proper authentication is complicated by factors such as the federation of virtual networks or mobility of virtual routers and links. Approaches that aim to deal with such difficulties are explained next.

5.2.1 Interoperability between federated virtual networks

Although isolation is one of the main security requirements in virtual networking, there are cases in which distinct virtual networks must be able to cooperate. The federation of virtual networks can, for example, enable end-to-end connectivity – through virtual devices of distinct virtual networks – or allow access to distinct services. However, it may not be possible to provide interoperability due to the heterogeneous nature of virtual networks (which may implement different, incompatible protocols). Chowdhury *et al.* [25] partially tackle this issue with a framework that manages identities in this kind of environment. The main objective of the work is to provide a global identification system. To this end, the authors employ a decentralized approach in which controllers and adapters are placed in each virtual network. Controllers provide functionalities such as address allocation and name resolution, while adapters act as gateways between virtual networks, performing address and protocol translations. The proposed global identification system does not restrict the internal identification mechanisms used locally by virtual networks, allowing each virtual network to keep its own internal naming scheme. Additionally, global identifiers used by this framework are unique, immutable, and not associated with physical location, in order to not hinder the security or mobility of virtual devices.

5.2.2 Certificate-based

As previously mentioned, the framework presented by Cabuk *et al.* [5] makes use of Trusted Virtual Domains (TVDs) to provide access control and network isolation. The authentication necessary to support access control is provided by means of digital certificates. These certificates ensure the identity of entities joining the network. Additionally, the system makes use of Virtual Private Networks (VPNs) to authenticate entities in network communications.

Analogously, Wolinsky *et al.* [19] use IPSec with X.509-based authentication for the purpose of access control in their system. In order to access the system, joining machines must request a certificate to the Certification Authority (CA). The CA responds by sending back a signed certificate to the node. The IP address of the requesting node is embedded into the certificate in order to prevent other nodes from reusing it.

5.2.3 Key-based

Fernandes and Duarte [26,31] present an architecture that aims to provide efficient routing, proper resource isolation and a secure communication channel between routers and the Virtual Machine Monitor (VMM) in a physical router. In order to ensure efficiency, virtual routers copy routing-related information to the VMM – in this case, the hypervisor. This process is performed by a plane separation module, which separates the data plane (which contains routing rules) and the control plane (responsible for creating routing rules). As a result, packets matching rules in the hypervisor routing table do not need to be redirected to virtual routers, resulting in a significant performance speedup. However, the process of copying routing information needs to be authenticated such that a malicious router is not able to compromise the data plane of another router.

In order to prevent identity fraud, the system requires mutual authentication between virtual routers and the VMM. Figure 4 depicts a simplified representation of the proposed architecture. The authors consider a Xen (paravirtualization)-based environment, in which virtual routers reside in unprivileged domains (DomUs) while the hypervisor resides within the privileged domain (Dom0). Each virtual router, upon instantiation, connects to the hypervisor following the client–server paradigm and performs an initial exchange of session keys using asymmetrical cryptography. The use of unique keys allows the hypervisor to verify the identity of distinct virtual routers in different unprivileged domains (in this example, DomU_1, DomU_2, and DomU_3) and to isolate traffic between them. After this initial key exchange, the secure

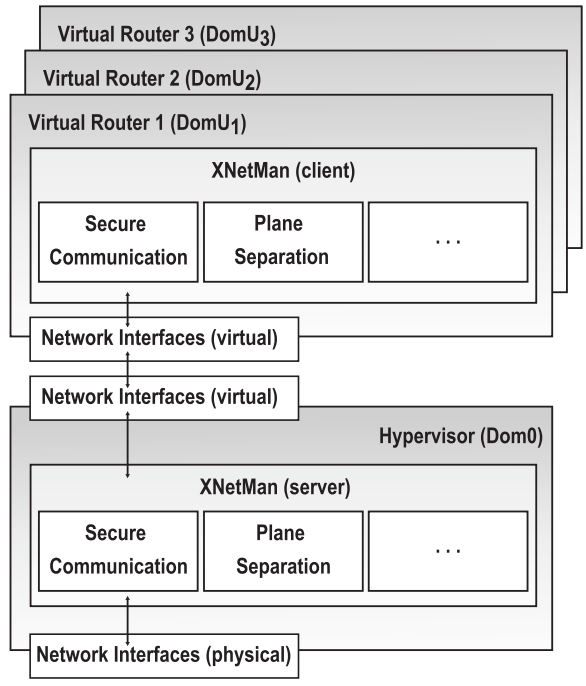

Figure 4 Simplified version of the architecture presented by Fernandes *et al.* [26,31], showing the secure communication modules.

communication module is used by other system modules in order to securely exchange messages with the hypervisor.

5.3 Data confidentiality

As network virtualization promotes the sharing of network devices and links among multiple entities, data confidentiality is a major security-related concern. Next, we explore approaches that leverage different protocols and techniques in order to provide secure communication within virtual networks.

5.3.1 VLANs and VPNs

The security goals approached by Cabuk *et al.* [5] include integrity, data isolation, confidentiality, and information flow control. Other than integrity, the remaining three goals, are directly related, and are tackled by a data confidentiality mechanism. The framework uses TVDs to control data access. However, virtual machines that belong to different TVDs may be hosted in the same physical machine. Therefore, it is necessary to ensure proper isolation, preventing a TVD from accessing data that belongs to another TVD.

The proposed solution for this challenge employs a combination of VLANs and VPNs. VLANs are used to identify packets belonging to different networks, allowing VLAN-enabled devices to route packets to the appropriate network interfaces, thus providing adequate isolation. Untrusted physical channels, however, may require a higher level of security. Therefore, if necessary, VPNs are used to provide data confidentiality by means of end-to-end cryptography.

5.3.2 Tunneling and cryptography

Wolinsky *et al.* [19] make use of tunneling in order to isolate network traffic between virtual machines (in this case, virtual workstations). Two tunneling approaches are employed. In the first approach, the host system runs a tunneling software that captures packets incoming from physical interfaces and forwards them to virtual machines. In the second approach, the tunneling software runs inside virtual machines, and traffic is restricted within virtual networks through the use of firewall rules. According to the authors, while the second approach is easier to deploy, malicious users may be able to subvert this firewall, compromising the system. Although the focus of Wolinsky *et al.* is isolation between virtual workstations, we believe that the techniques used to achieve such isolation could be extended to virtual routers in network virtualization environments.

Fernandes and Duarte [26,31] deal with data confidentiality in communications between a virtual router and the Virtual Machine Monitor (VMM) hosting it. After the authentication process, described in Section 5.2, virtual routers use symmetrical cryptography in order to securely communicate with the VMM.

Huang *et al.* [22] present a framework that provides secure routing. In the environment presented by the authors, routing information that is propagated through a virtual network is confidential and needs to be kept secret from unauthorized network entities. Routing information is categorized in groups, and group keys are assigned to virtual routers. Therefore, routing information can be encrypted, ensuring that only routers with the correct key are able to decrypt this information. Thus, routing information relative to a given group is protected against unauthorized access from other groups, other virtual networks or the physical network itself.

Similarly to the previously described approach, Fukushima *et al.* [24] aim to protect sensitive routing information in virtual networks from being disclosed to entities controlling the physical network. To achieve this goal, the authors make use of a strategy based on Secure Multi-party Computation (SMC). SMC allows multiple entities to perform joint computations on sensitive data they hold without disclosing such data. Each entity has access to the result of the global computation, but not to any data held by other entities. This is achieved through the use of one-way functions, which are easy to evaluate but hard to invert. In the context of virtual network routing, SMC allows a virtual router to compute optimal

routes without needing to share the information that it holds. As SMC requires full-mesh connectivity between computing nodes, the authors decompose the virtual network into locally connected subsets of routers, called *cliques*. The SMC-based distributed routing algorithm is run locally in each *clique*, and the results of local computations are then shared between *cliques*.

As the employment of cryptographic techniques requires physical devices that are capable of supporting protocols that enable them and generates processing and bandwidth overheads, Bays *et al.* [4] devise an optimization model and a heuristic algorithm for online, privacy-oriented virtual network embedding. Clients may require end-to-end or point-to-point cryptography for their networks, as well as requiring that none of their resources overlap with other specific virtual networks. Both the optimal and heuristic approaches take into account whether physical routers are capable of supporting cryptographic algorithms in order to ensure the desired level of confidentiality and guarantee the non-overlapping of resources (if requested). Additionally, both methods feature precise modeling of overhead costs of security mechanisms in order to not underestimate the capacity requirements of virtual network requests. This proposal is in line with research performed in the area of virtual network embedding, such as the work of Alkmim *et al.* [41].

5.3.3 Firewalling and subnetting

As previously mentioned in Section 5.3.2, Wolinsky *et al.* [19] make use of firewall rules (in addition to tunneling techniques) in order to prevent communications between different virtual networks. In addition to using firewalls for this purpose, Wu *et al.* [20] also employ subnetting (*i.e.*, each virtual network is bound to a unique subnet) in order to provide an additional layer of security against unauthorized information disclosure.

5.3.4 Path splitting

In addition to encryption of routing information, Huang *et al.* [22] use variable paths in virtual networks to propagate data flows. Figure 5 illustrates the employment of path splitting in order to hinder an information interception attack. Communication between a virtual router hosted on Physical Router (PR) 1 and another one hosted on PR 7 is split among two different paths – one passing through PR 3 and 6, and the other, through PR 2 and 4 (represented by dashed lines). Even if traffic between these two virtual routers is not encrypted, the threat is partially mitigated as the attacker only has access to part of the information being exchanged (packets passing through the link between PR 3 and 6). Moreover, when used in combination with encryption (as in the work of Huang *et al.*), this approach helps mitigate traffic analysis attacks.

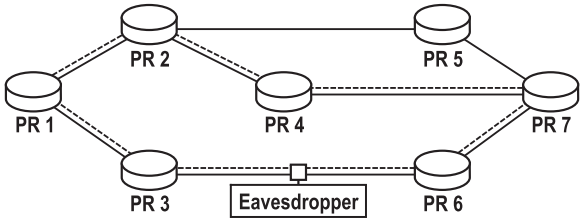

Figure 5 Example of path splitting used to mitigate an information interception attack.

It is worth noting that while in this example the attacker is only eavesdropping on one physical path, in reality, multiple devices may be compromised. In this case, splitting traffic among an increasing number of paths would lead to progressively higher levels of security (or, conversely, to increasingly higher costs for an attacker to capture the full traffic).

5.3.5 Limiting introspection

Finally, van Cleeff *et al.* [10] present recommendations for safer use of virtualization. One of these recommendations is to limit, or even disable, the introspection feature, which allows virtual machine monitors to access data inside virtual machines. While useful, this functionality may be exploited by attackers, as previously explained on Subsection 4.1.3.

5.4 Data integrity

Similarly to confidentiality, data integrity is a major concern as a result of shared network devices and communication channels. Next, we describe approaches that aim to establish a desired level of integrity in virtual network environments.

5.4.1 Cryptography

In addition to authentication (*i.e.*, source integrity) and confidentiality, the framework developed by Cabuk *et al.* [5] makes use of VPNs to provide data integrity to virtual networks. The use of cryptographic tunneling protocols prevents malicious entities from manipulating messages going through the network. As previously discussed, the authors use IPSec as the tunneling protocol.

5.4.2 Timestamping

As previously discussed, replay attacks are one of the threats to data integrity that may be present in network virtualization environments. The addition of unique identifiers inside encrypted messages makes it possible to detect duplicated messages, and therefore, replay attacks. For this purpose, the architecture proposed by Fernandes and Duarte [26,31] inserts timestamps inside encrypted messages in order to ensure that messages are non-reproducible.

5.4.3 Limiting introspection

Besides mitigating information theft, disabling or limiting introspection also prevents data tampering. According to van Cleeff *et al.* [10], this functionality allows the VMM to modify applications running inside it, which may cause inconsistencies. Another recommendation consists of specifically designing applications that facilitate batch processing and checkpointing. According to the authors, this minimizes security issues associated with rollback and restore operations that may otherwise threaten integrity.

5.5 Nonrepudiation

Nonrepudiation provides evidences regarding which (potentially malicious) actions have been performed by which entities. This security countermeasure is highly valuable in the context of network virtualization environments, in which a number of physical devices are shared by different users. Nevertheless, we are not aware of any publication that targets this countermeasure specifically.

5.6 Availability

Last, we present proposals that aim to maintain the availability of network virtualization environments. The key concerns in this area of security are providing proper resource isolation and mitigating attacks that target physical or virtual devices. Approaches aiming to deal with such concerns are explained in the following subsections.

5.6.1 Physical resource isolation

One of the main concerns regarding availability is the abuse of physical resources by virtual networks. Virtual networks may attempt to use as much resources as possible in order to maximize their performance. If the environment is not adequately protected, this behavior may lead to the exhaustion of physical resources, compromising the availability of other virtual networks hosted on the same substrate. Therefore, physical resources must be shared in a fair manner, and actions performed by a virtual network must not negatively impact others.

According to Wu *et al.* [29], the sharing of physical resources by packet processors is usually only performed at a granularity of entire processor cores. The authors claim that finer-grained processor sharing is required in order to provide scalability for network virtualization environments. Thus, the authors propose a system that allows multiple threads to share processor cores concurrently while maintaining isolation and fair resource sharing. However, typical multithreading approaches consider a cooperative environment, which is not the case in network virtualization. The authors devise a fair multithreading mechanism that allows the assignment of different priorities to each thread. Additionally, this mechanism takes into account the history of how much processing has been performed by each thread. Inactivity times are also considered in order to guarantee that threads will not stay idle for too long. The evaluation performed by the authors shows that the proposed mechanism is able to properly distribute processing resources according to the defined priorities. Furthermore, while it requires more processing power, it is able to provide better resource utilization in comparison to coarse-grained approaches.

Kokku *et al.* [30] propose a network virtualization scheme that provides resource isolation while aiming to maximize substrate utilization. It allows virtual networks to have either resource-based reservations (*i.e.*, reservations calculated as a percentage of available resources in the substrate) or bandwidth-based reservations (*i.e.*, reservations based on the aggregate throughput of the virtual network). Virtual networks are divided in two groups according to the type of reservation required, and treated independently by a scheduler. This scheduler treats flows that belong to different virtual networks with distinct priorities, based on the reservations and average resource usage rate of each network. The authors present an evaluation performed on an implemented prototype, showing that the proposed scheme was capable of ensuring that each virtual network met its reservations.

Fernandes and Duarte [26] present a network monitor that employs plane separation in order to provide resource isolation in network virtualization environments. The system is able to allocate resources based on fixed reservations, as well as to redistribute idle resources between virtual networks that have a higher demand. Additionally, an administrator is able to control the amount of resources to be used by each virtual network, as well as set priorities for using idle resources. The system continuously monitors the consumption of physical resources by each virtual router. If any virtual router exceeds its allowed use of bandwidth, processing power, or memory, it is adequately punished by having packets dropped, or a percentage of its stored routes erased. Harsher punishments are instituted if there are no idle resources available. Conversely, given punishments are gradually reduced if the router stops using more than its allocated resources. This system is capable of adequately preventing physical resources from being overloaded, and packet drops employed by the punishment mechanism do not cause a major impact on network traffic.

In another publication [31], the same authors extend the previously described network monitor. This new system introduces the idea of short term and long term requirements, based on the time frame in which they must be met. Short term requirements may be allocated in an exclusive or non-exclusive manner, while long term requirements are always non-exclusive. In this context, exclusive requirements are always allocated (even if part of the allocated resources is idle), while non-exclusive

requirements are only allocated when necessary. The system prioritizes virtual networks that have used the lowest portion of their requirements, and an adaptive control scheme is used in order to improve the probability that long term requirements, if needed, will be met. The presented evaluation shows the improvement of this system over the original [26] in terms of guaranteeing that the demands of each virtual network will be met, as well as reducing resource load on the physical substrate.

5.6.2 Virtual network resilience

Even with proper physical resource isolation, maintaining availability remains a challenge in virtualized networks. The virtualization layer must be resilient, maintaining its performance and mitigating attacks in order to sustain its availability. Some of the publications described next approach the issue of virtual network resilience from the point of view of fault tolerance. Nonetheless, we emphasize that the solutions described in these publications may also be used as a response to attacks that cause the failure or degradation of physical devices or links.

The solution presented by Yeow *et al.* [32] aims to provide network infrastructures that are resilient to physical router failures. This objective is achieved through the use of backups (*i.e.,* redundant routers and links). However, redundant resources remain idle, reducing the utilization of the physical substrate. To minimize this problem, the authors propose a scheme that dynamically creates and manages shared backup resources. This mechanism minimizes the number of necessary backup instances needed to achieve a certain level of reliability. While backup resources are shared, each physical router is restricted to hosting a maximum number of backup instances in order not to sacrifice reliability. The connectivity between each virtual router and its neighbors is preserved in all of its backups, both in terms of number of links and bandwidth reservations.

The illustration on the left side of Figure 6 shows a simple representation of how backup nodes (represented as circles) may be shared among different virtual networks. For example, the two backup nodes at the right side of this figure are shared between Virtual Network 1 and Virtual Network 3, regardless of whether they belong to one or the other. The right side of Figure 6, in turn, depicts in greater detail how backups are allocated to virtual routers. A virtual router C_1 has virtual routers B_1 and B_2 as its backups. Since C_1 has a virtual link connecting it to another router, N_1, a virtual link with the same bandwidth reservation (depicted as 1 in the figure) is also established between each backup node and N_1 in order to preserve the connectivity of the original router.

Meixner *et al.* [35] devise a probabilistic model for providing virtual networks that are resilient to physical disasters. Disasters are characterized by the occurrence of multiple failures in the physical network, as well as the possibility of correlated cascading failures during attempts to recover network resources. The virtual link mapping strategy guarantees that the failure of a single physical link will not disconnect any virtual network, and aims at minimizing virtual network disconnection in the event of a disaster (*i.e.*, simultaneous failure of multiple links). Additionally, excess processing capacity in the physical network is used to create a backup router for each virtual network, which reduces disconnection in the event of disasters and provides additional processing capacity for the recovery phase. When attempting to recover virtual network resources, the model analyzes all possible virtual router replacements in an effort to replace affected virtual routers in a way that ensures the virtual network will not be disconnected by any post-disaster failures.

The system presented by Zhang *et al.* [27] uses redundant virtual networks in order to provide reliable live streaming services. It is able to detect path failures and traffic congestion, dynamically redirecting data flows.

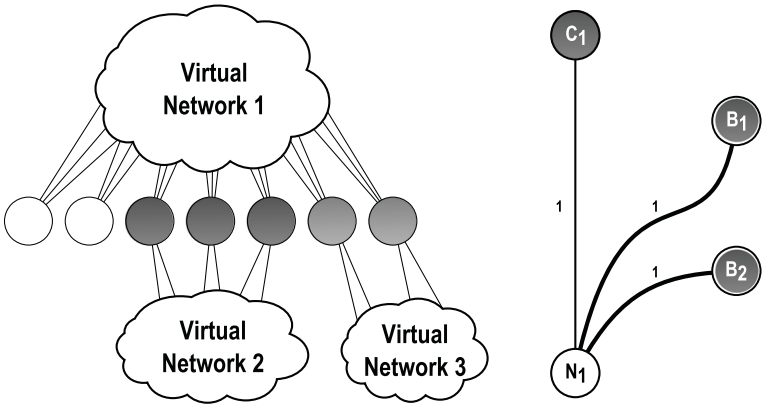

Figure 6 Examples of sharing and mapping of backup instances, used by Yeow *et al.* [32] to provide resilient virtual networks.

Initially, the data flow is distributed equally through available virtual networks. Figure 7 depicts the distribution of a data flow through virtual networks, using multiple paths between a server and a client. Gradually, the number of packets routed through each virtual network is adapted according to its relative bandwidth capacity. Additionally, an active probing mechanism is used to detect failures in the physical network or routing problems (changes in routing tables, for example, may have a significant impact in live streaming applications). If an issue is detected, the system is able to redirect data flows away from problematic networks and redistribute it among the remaining ones. Experiments performed by the authors demonstrate advantages in using multiple networks instead of a single one, with increasing gains when using up to four virtual networks. Additionally, the authors claim that the bandwidth cost of the probing mechanism is negligible.

Chen *et al.* [33] propose a virtual network embedding strategy that aims at ensuring survivability. Load balancing is employed in the embedding process in order to balance the bandwidth consumption of substrate links. Moreover, backup links are reserved for each accepted virtual network, but not activated until a failure occurs. Backup links are allocated in physical paths that do not overlap with the path hosting the original link, guaranteeing that a single physical link failure will not simultaneously affect the original virtual link and one or more of its backups. These backup resources may be shared by multiple virtual networks or reconfigured over time in order to improve efficiency.

Zhang *et al.* [34] devise a strategy for computing the availability of Virtual Data Centers (VDCs), as well as an algorithm for reliable VDC embedding. In order to determine VDC availability, the authors consider the availability of individual, heterogeneous components, as well as dependencies among them. The embedding mechanism aims at meeting minimum availability criteria while optimizing resource usage. Virtual devices are divided into replication groups (groups in which any virtual device may serve as a backup if another fails). In order to minimize resource consumption, VDCs are embedded on physical devices with the lowest level of availability that still meets the desired level. In a similar way, a minimum number of backups is assigned to each replication group in order to meet availability requirements.

Unlike the previously described approaches in the area of virtual network resilience, Oliveira *et al.* [7] present a strategy based on "opportunistic resilience", which does not employ backup resources. The bandwidth demand of each virtual link is split over multiple physical paths. As a consequence, physical link failures are less likely to cause a virtual link disconnection (an affected virtual link will remain operational, albeit with less capacity). Additionally, when link failures occur, a reactive strategy is used in order to reallocate the lost capacity over unaffected paths, attempting to fully restore the bandwidth of degraded virtual links.

Distributed Denial of Service (DDoS) attacks are a common threat to the availability of network services. The system proposed by Yu and Zhou [6] aims to detect such attacks on community networks (federated virtual networks that belong to cooperating entities). The devised solution leverages communication between virtual routers that belong to different entities in this collaborative environment to detect possible attacks at an early stage. Virtual routers located on the edges of the community network monitor traffic passing through them and calculate the entropy of its flows. Traffic surges in any of these flows will cause the entropy to drop, indicating a possible attack. If this occurs, other routers are notified

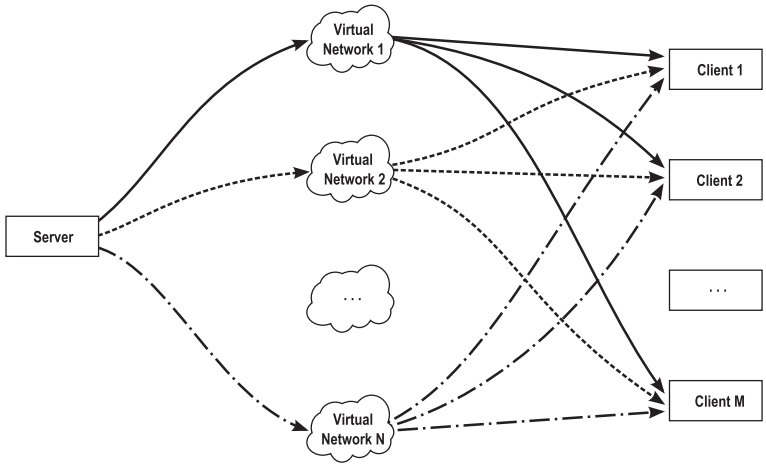

Figure 7 A live streaming data flow is distributed among different virtual networks, a mechanism used by Zhang *et al.* [27].

and instructed to calculate the entropy rate of this suspected flow. Calculated values are compared, and if they are similar, a DDoS attack is confirmed.

6 Discussion

A number of insights can be obtained from the extensive investigation of the state of the art reported in this paper. First, it is possible to observe that the publications in the area are not equally distributed between the main security categories. Tables 3 and 4 show, respectively, the security threats and security countermeasures approached in these publications. In both tables, publications have been grouped together according to the security elements they approach, whenever possible. It is noticeable that disruption and availability – a security threat and a countermeasure that are directly correlated – are approached in the majority of these publications. This is likely due to the high prevalence of attacks aiming at causing disruption. These attacks are relatively simple but can be highly devastating, especially in an environment that makes heavy use of shared resources (as a single physical failure may disrupt several virtual networks). Disclosure and confidentiality follow closely behind, being present in a similar number of publications as disruption/availability. Once again, this is linked to physical resource sharing. Similarly to disruption attacks, such sharing means that a single well-placed sniffer may be able to acquire sensitive information from multiple virtual networks at once. Moreover, there are also privacy concerns between infrastructure providers and virtual network requesters (as the former may have access to data that the latter considers confidential).

Second, only a small number of publications approach more than one threat or countermeasure simultaneously. No single publication has dealt with threats in more than two of the four categories, or presented solutions that provide more than four security countermeasures, out of a total of six. Additionally, one security countermeasure in particular – nonrepudiation – was not approached by any of the publications. The combination of authentication and integrity, which exists in some publications, can be considered as the basis for the provision

Table 3 Security threats mentioned in publications in the area of virtual network security

Publication	Threats			
	DI	DE	DR	US
[4]	×			
[19]	×			
[21]	×			
[22]	×			
[23]	×			
[24]	×			
[10]	×	×		
[20]	×	×		
[5]	×	×		×
[25]		×		
[26]		×	×	
[36]				×
[6]			×	
[7]			×	
[27]			×	
[28]			×	
[29]			×	
[30]			×	
[31]			×	
[32]			×	
[33]			×	
[34]			×	
[35]			×	

From left to right: Disclosure, Deception, Disruption, Usurpation.

Table 4 Security countermeasures provided by publications in the area of virtual network security

Publication	Countermeasures					
	AC	AU	CO	IN	NR	AV
[4]			×			
[20]			×			
[22]			×			
[24]			×			
[19]	×	×	×			
[5]	×	×	×	×		
[26]		×	×	×		×
[31]		×	×	×		×
[10]			×	×		
[25]		×				
[6]						×
[7]						×
[27]						×
[29]						×
[30]						×
[32]						×
[33]						×
[34]						×
[35]						×

From left to right: Access Control, Authentication, Confidentiality, Integrity, Nonrepudiation, Availability.

of nonrepudiation, but this specific countermeasure is not targeted. Nonrepudiation is a highly valuable (albeit challenging) security countermeasure for network virtualization environments, and will be further discussed in Section 7.

Third, we were able to conclude that many of the threats that affect network virtualization environments also affect traditional networks. However, we emphasize that these threats affect traditional and virtual network environments in different ways. In most cases, the effects of these threats are greatly exacerbated by certain characteristics of virtual network environments. Information interception, physical resource overloading, physical resource failure, and software vulnerability exploitation are aggravated by the fact that a number of virtual routers may share a physical router. Therefore, as previously explained, an attack of any of these types targeting a single physical router may affect several virtual networks. Further, it is more difficult to recognize (and, therefore, to perform countermeasures against) identity fraud and replay attacks due to the dynamicity of network virtualization environments, as virtual routers may be freely moved among physical routers and assume different identities. Loss of registry entries and information leakage, as described in the studied literature, are limited to virtual network environments. Moreover, threats related to introspection are also inherent to these types of environments, as this is a (potentially exploitable) feature of virtual machine monitors.

Last, we can observe the employment of different virtualization techniques in some publications. For example, Cabuk *et al.* [5] implemented a prototype of their framework based on a paravirtualization platform, while Huang *et al.* [22] consider an underlying network based on programmable routers. Further, Fernandes and Duarte [26,31] build a hybrid solution that combines paravirtualization with plane separation, a core idea of programmable networks. Although the majority of publications do not target specific network virtualization techniques, we emphasize that different types of platforms have their own sets of benefits, as well as security concerns, which need to be taken into account.

7 Challenges

Despite the existence of a sizable body of work in virtual network security, some challenges remain open. In this section, we summarize some of the main research challenges in this area. We emphasize, however, that these challenges should not be considered exhaustive, but rather as a starting point for further discussions in the area.

One clear opportunity for research in virtual network security is the provisioning of nonrepudiation – which, to the best of our knowledge, has not yet been approached. Nonrepudiation requires providing proof of actions performed by entities on a network, which can be used for holding entities accountable for malicious activity. We deem nonrepudiation an essential security countermeasure for virtual networking environments in order to accurately backtrace attacks – not only to ensure that punitive actions will be taken against the attackers but also to properly contain the attacks themselves. In the event of a DDoS attack, for example, this countermeasure could enable administrators to pinpoint the origins of the attack with a high level of precision – which otherwise tends to be a very difficult task. Moreover, nonrepudiation may even prevent attacks, in the sense that malicious users who are aware that such a mechanism is in use may refrain from carrying out attacks in order to avoid exposing themselves. Provisioning nonrepudiation can be challenging for a number of reasons, such as the complexity of securely storing and handling digital certificates – used for proving that an action was, indeed, performed by a given entity – and the negative impact this has on network performance. Moreover, it is necessary to maintain a desired level of privacy for virtual network requesters as well as end users. Nevertheless, we envision that the importance of this countermeasure will grow steadily as network virtualization becomes increasingly prominent in production environments.

In addition to privacy issues related to nonrepudiation, there are also concerns regarding the privacy of general data stored in virtual routers or sent through virtual networks. Although such data may be protected from being intercepted by other entities, infrastructure providers have physical access to all data stored in virtual networks they are hosting. Although this issue has been approached by some authors, their proposed strategies are often based on strong assumptions, such as the ability to choose which physical entity (out of a number of entities controlling the physical substrate) will host each of its routers – a feature that may not commonly be available in practice.

Another opportunity stems from the multiple levels of heterogeneity present in network infrastructures. As previously mentioned, in addition to the use of heterogeneous hardware devices, it is common for network substrates to be composed of a number of physical networks that belong to different entities. As such, there is a need for uniform methods for requesting, negotiating, and enforcing security requirements across devices that may have incompatible interfaces and entities with potentially conflicting policies.

Last, software platforms used to instantiate virtual networks may not always offer adequate protection against security threats. Moreover, although virtualization technologies are gradually evolving and becoming more mature, both hardware and software are susceptible to vulnerabilities that may be exploited by attackers.

Consequently, research efforts that build on top of network virtualization need to consider these security issues and, most importantly, overhead costs of additional security mechanisms that may be necessary, in order to ensure that they will be suitable for real world environments.

We emphasize, once again, that this is not an exhaustive list of challenges in the area. The essence of network virtualization is based on layers upon layers with increasing levels of abstraction (*e.g.*, the physical substrate, the virtualization layer, virtual networks, and services running on top of them). Consequently, we envision that a number of other challenges may be present in all of these layers – much like the ones listed in this section.

8 Conclusions

Network virtualization enables the subdivision of a single network infrastructure into multiple virtual architectures. The benefits of this technique apply to a wide range of applications, including the creation of virtual testbeds, community networks, and cloud computing infrastructures. Furthermore, network virtualization has been proposed by researchers as the basis for the creation of a new architecture for the Internet, allowing pluralist network environments that support a number of different network protocols simultaneously.

In spite of the benefits provided by network virtualization, there is a series of security issues that need to be considered. Our study revealed a number of security threats, covering the four categories defined by Shirey [16]. The very act of sharing a physical infrastructure among multiple parties is shown to be the source of several of these threats.

This study shows that there have been several efforts to provide security in virtual networks. However, these efforts were not organized in a comprehensible manner. This study provides a systematic overview of the available research results in the field, categorizing work that represents the state of the art and highlighting different approaches for providing security. Additionally, it also evidences imbalances between different sub-areas of security research in network virtualization, which can be used as guidance for future work in this area. Usurpation and access control, for example, are significantly underrepresented in relation to other security countermeasures, and nonrepudiation is not targeted by any publication. Additionally, while a significant body of work exists in the sub-area of availability, only one publication deals with detection and prevention of attacks. Such gaps may represent valuable opportunities for future work.

To summarize, the categorization of security threats and countermeasures presented in this paper simplifies the analysis of which security aspects have not yet been approached and which types of threats need to be mitigated. Furthermore, it makes it easier to identify a number of existing solutions that aim to provide security in virtual networks.

Endnotes

[1]Machine virtualization is available for personal computers, in commonly used operating systems (e.g., Windows, Linux, and Mac OS X).

Competing interests
The authors declare that they have no competing interests.

Authors' contributions
All authors read and approved the final manuscript.

Acknowledgements
The authors would like to thank the anonymous reviewers for their valuable comments and suggestions. This work has been partially supported by FP7/CNPq (Project SecFuNet, FP7-ICT-2011-EU-Brazil), RNP-CTIC (Project ReVir), as well as PRONEM/FAPERGS/CNPq (Project NPRV).

Author details
[1]Institute of Informatics, Federal University of Rio Grande do Sul, Porto Alegre, Brazil. [2]Institute of Computing, University of Campinas, Campinas, Brazil.

Received: 29 August 2014 Accepted: 8 December 2014
Published online: 27 January 2015

References

1. Chowdhury NMMK, Boutaba R (2010) A survey of network virtualization. Comput Netw 54(5):862–876
2. Fernandes N, Moreira MD, Moraes I, Ferraz L, Couto R, Carvalho HT, Campista M, Costa LK, Duarte OB (2011) Virtual networks: isolation, performance, and trends. Ann Telecommun 66(5–6):339–355
3. Anderson T, Peterson L, Shenker S, Turner J (2005) Overcoming the internet impasse through virtualization. Computer 38(4):34–41
4. Bays LR, Oliveira RR, Buriol LS, Barcellos MP, Gaspary LP (2014) A heuristic-based algorithm for privacy-oriented virtual network embedding. In: IEEE/IFIP Network Operations and Management Symposium (NOMS). IEEE, Krakow, Poland
5. Cabuk S, Dalton CI, Ramasamy H, Schunter M (2007) Towards automated provisioning of secure virtualized networks. In: ACM Conference on Computer and Communications Security. New York, USA
6. Yu S, Zhou W (2008) Entropy-based collaborative detection of ddos attacks on community networks. In: IEEE International Conference on Pervasive Computing and Communications. IEEE Computer Society, Washington, DC, USA
7. Oliveira RR, Marcon DS, Bays LR, Neves MC, Buriol LS, Gaspary LP, Barcellos MP (2013) No more backups: Toward efficient embedding of survivable virtual networks. In: IEEE International Conference on Communications. IEEE, Budapest, Hungary
8. LAN/MAN Standards Committee (2006) IEEE Standard for Local and metropolitan area networks – Virtual Bridged Local Area Networks. IEEE Std 802.1Q-2005 (incorporates IEEE Std 802.1Q1998, IEEE Std 802.1u-2001, IEEE Std 802.1v-2001, and IEEE Std 802.1s-2002). http://www.ieee802.org/1/pages/802.1Q-2005.html
9. Rosen E, Cisco Systems I, Rekhter Y, Juniper Networks I (2006) RFC 4364: BGP/MPLS IP Virtual Private Networks (VPNs). http://www.ietf.org/rfc/rfc4364.txt
10. van Cleeff A, Pieters W, Wieringa RJ (2009) Security implications of virtualization: A literature study. In: International Conference on Computational Science and Engineering. IEEE Computer Society, Washington, DC, USA
11. McKeown N, Anderson T, Balakrishnan H, Parulkar G, Peterson L, Rexford J, Shenker S, Turner J (2008) Openflow: enabling innovation in campus networks. SIGCOMM Comput Commun Rev 38:69–74
12. Bari MF, Boutaba R, Esteves R, Granville LZ, Podlesny M, Rabbani MG, Zhang Q, Zhani MF (2013) Data center network virtualization: A survey. Communications Surveys Tutorials, IEEE 15:909–928

13. Zhou M, Zhang R, Xie W, Qian W, Zhou A (2010) Security and privacy in cloud computing: A survey. In: Semantics Knowledge and Grid (SKG), 2010 Sixth International Conference On. IEEE, Beijing, China
14. Hashizume K, Rosado DG, Fernández-Medina E, Fernandez EB (2013) An analysis of security issues for cloud computing. J Internet Serv Appl 4:1–13
15. Scott-Hayward S, O'Callaghan G, Sezer S (2013) Sdn security: A survey. In: Future Networks and Services (SDN4FNS), 2013 IEEE SDN For. IEEE, Trento, Italy
16. Shirey R (2000) RFC 2828: Internet Security Glossary. http://www.ietf.org/rfc/rfc2828.txt
17. Stallings W (2006) Cryptography and Network Security: Principles and Practice. Pearson/Prentice Hall, Upper Saddle River, New Jersey, USA
18. Cavalcanti E, Assis L, Gaudencio M, Cirne W, Brasileiro F (2006) Sandboxing for a free-to-join grid with support for secure site-wide storage area. In: International Workshop on Virtualization Technology in Distributed Computing. IEEE Computer Society, Washington, USA
19. Wolinsky DI, Agrawal A, Boykin PO, Davis JR, Ganguly A, Paramygin V, Sheng YP, Figueiredo RJ (2006) On the design of virtual machine sandboxes for distributed computing in wide-area overlays of virtual workstations. In: International Workshop on Virtualization Technology in Distributed Computing. IEEE Computer Society, Washington, DC, USA
20. Wu H, Ding Y, Winer C, Yao L (2010) Network security for virtual machine in cloud computing. In: Computer Sciences and Convergence Information Technology (ICCIT), 2010 5th International Conference On. IEEE, Seoul, South Korea
21. Cui Q, Shi W, Wang Y (2009) Design and implementation of a network supporting environment for virtual experimental platforms. In: WRI International Conference on Communications and Mobile Computing. IEEE Computer Society, Washington, DC, USA
22. Huang D, Ata S, Medhi D (2010) Establishing secure virtual trust routing and provisioning domains for future internet. In: IEEE Conference on Global Telecommunications, Miami, USA
23. Pignolet Y-A, Schmid S, Tredan G (2013) Adversarial vnet embeddings: A threat for isps?. In: IEEE INFOCOM. IEEE, Turin, Italy
24. Fukushima M, Sugiyama K, Hasegawa T, Hasegawa T, Nakao A (2013) Minimum disclosure routing for network virtualization and its experimental evaluation. IEEE/ACM Trans Netw PP(99):1839–1851
25. Chowdhury NMMK, Zaheer F-E, Boutaba R (2009) imark: an identity management framework for network virtualization environment. In: IFIP/IEEE International Symposium on Integrated Network Management. IEEE Press, Piscataway, USA
26. Fernandes NC, Duarte OCMB (2011) Xnetmon: A network monitor for securing virtual networks. In: IEEE International Conference on Communications. IEEE, Kyoto, Japan
27. Zhang Y, Gao L, Wang C (2009) Multinet: multiple virtual networks for a reliable live streaming service. In: IEEE Conference on Global Telecommunications. IEEE Press, Piscataway, USA
28. Marquezan CC, Granville LZ, Nunzi G, Brunner M (2010) Distributed autonomic resource management for network virtualization. In: IEEE/IFIP Network Operations and Management Symposium, Osaka, Japan
29. Wu Q, Shanbhag S, Wolf T (2010) Fair multithreading on packet processors for scalable network virtualization. In: ACM/IEEE Symposium on Architectures for Networking and Communications Systems. ACM, New York, USA
30. Kokku R, Mahindra R, Zhang H, Rangarajan S (2010) Nvs: a virtualization substrate for wimax networks. In: International Conference on Mobile Computing and Networking. ACM, New York, USA
31. Fernandes NC, Duarte OCMB (2011) Provendo isolamento e qualidade de serviço em redes virtuais. In: Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos, Campo Grande, Brazil. (in Portuguese)
32. Yeow W-L, Westphal C, Kozat UC (2011) Designing and embedding reliable virtual infrastructures. SIGCOMM Comput Commun Rev 41(2):57–64
33. Chen Q, Wan Y, Qiu X, Li W, Xiao A (2014) A survivable virtual network embedding scheme based on load balancing and reconfiguration. In: IEEE Network Operations and Management Symposium. IEEE, Krakow, Poland
34. Zhang Q, Zhani MF, Jabri M, Boutaba R (2014) Venice: Reliable virtual data center embedding in clouds. In: IEEE INFOCOM. IEEE, Toronto, Canada
35. Meixner CC, Dikbiyik F, Tornatore M, Chuah C, Mukherjee B (2013) Disaster-resilient virtual-network mapping and adaptation in optical networks. In: International Conference on Optical Network Design and Modeling
36. Roschke S, Cheng F, Meinel C (2009) Intrusion detection in the cloud. In: IEEE International Conference on Dependable, Autonomic and Secure Computing. IEEE Computer Society, Washington, DC, USA
37. Common Vulnerabilities and Exposures (2012) CVE-2012-1516. https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2012-1516
38. Common Vulnerabilities and Exposures (2012) CVE-2012-1517. https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2012-1517
39. Common Vulnerabilities and Exposures (2012) CVE-2012-2449. https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2012-2449
40. Common Vulnerabilities and Exposures (2012) CVE-2012-2450. https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2012-2450
41. Alkmim GP, Batista DM, Fonseca NLS (2013) Mapping virtual networks onto substrate networks. J Internet Serv Appl 3(4):1–15