Challenges for Cloud Networking Security

Peter Schoo[1], Volker Fusenig[1], Victor Souza[2], Márcio Melo[3], Paul Murray[4],
Hervé Debar[5], Houssem Medhioub[5] and Djamal Zeghlache[5]

[1] Fraunhofer Institute for Secure Information Technology SIT,
Garching near Munich, Germany
{peter.schoo,volker.fusenig}@sit.fraunhofer.de,
[2] Ericsson Research, Stockholm, Sweden
victor.souza@ericsson.com,
[3] Portugal Telecom Inovação, Aveiro, Portugal
marcio-d-melo@ptinovacao.pt
[4] HP Labs, Bristol, United Kingdom
pmurray@hp.com
[5] Institut Telecom, Telecom SudParis, France
{firstname.lastname}@it-sudparis.eu

Abstract. Cloud computing is widely considered as an attractive service model since the users commitments for investment and operations are minimised, and costs are in direct relation to usage and demand. However, when networking aspects for distributed clouds are considered, there is little support and the effort is often underestimated. The project SAIL is addressing *cloud networking* as the combination of management for cloud computing and vital networking capabilities between distributed cloud resources involved to improve the management of both. This position paper presents new security challenges as considered in SAIL for ensuring legitimate usage of cloud networking resources and for preventing misuse.

Key words: Cloud Networking, Cloud Computing, Network Virtualisation, Security

1 Introduction

Despite the fact that many applications are not ready to be deployed in a cloud computing environment the cloud is here to stay. Initially driven by the deployment of IT applications leveraging the economy of scale and multi-tenancy, the use of cloud computing has expanded to a much more diversified one. Applications that show high degree of variable demand for resources fit the cloud computing model well. The advantages of running applications in the cloud are manifold: lower costs through shared computing resources, no upfront infrastructure costs, and on-demand provisioning of computing nodes to fit transient requirements. Virtualisation in the data centres has been a key enabler to allow the dynamic provisioning of computing resources to become a reality.

An aspect that has been so far overlooked is the fact that the cloud computing model relies on connectivity end-to-end and thus on the network in between the user and the cloud. As applications move to the cloud, more will be demanded from existing networks in terms of capacity (likely more data to be sent across network links), quality (low delay for interactive applications), security, and many more.

Clearly, cloud applications will demand a network that is more flexible. Since applications and entire cluster of servers can be relocated to another data centre all of the networking pipes need to be *re-plumbed*. Consider a set of servers connected through a complex group of VLANs and L2 tunnels in a data centre. Moving one or many of those virtual machines will demand the network to be reconfigured in a matter of milliseconds. Networks that can swiftly be reconfigured will enable the full benefits of the cloud environment. This is the envisioned concept of cloud networking - it encompasses provisioning of on-demand network resources in a time span that is compatible with the allocation of computing resources in a cloud today.

Being network wise close to the user has its advantages as well. Network conditions such as latency may hinder the execution of certain cloud applications (e.g., virtual desktops) in certain locations. Thus we consider a cloud that is fully distributed through the network, with some servers closer to the user, others in the aggregation or access network, and others in data centres. Depending on the requested load one may need more servers in a certain geographical region. A geographically distributed cloud will enable finer control over the user experience.

Regardless, one of the top concerns of customers willing to deploy applications and store data in the cloud is security. Cloud networking should secure isolation of resources, provide misuse protection, confidentiality and integrity, authentication, authorisation, and auditing mechanisms.

This paper presents the research challenges of providing a secure cloud network system. These research challenges will be explored in the course of a 30 months project called SAIL (Scalable Adaptive Internet soLutions). SAIL [1] is an EU funded project whose consortium includes 24 partners from industry, academia, and research institutes. SAIL aims at creating technology to address some of the shortcomings of the current Internet which include the lack of a content-centric model for large scale content distribution; support for connectivity services providing point-to-multipoint capabilities; insufficient support for deployment of dynamic network slices in a cloud computing scenario.

Amongst others, SAIL will develop networking functions for applications with highly variable demands and integrate networking with cloud computing and storage, along with the necessary tools for management and security. In that way, the allocation of both computing and networking resources will be solved as only one optimisation problem. SAIL will provide a migration path whereby developed technologies will be deployed in the existing Internet.

Besides the cloud networking security related requirements and challenges, more fundamental cloud security aspects need to be addressed and considered. Cloud computing environments are likely to suffer from a number of known vulnerabilities, enabling attackers to either obtain computing services for free (attack against cloud providers), steal information from cloud users (attack against cloud customers data) or penetrate the infrastructure remaining in client premises through cloud connections (attack against cloud customer infrastructures). Typical examples of these attacks today are VoIP free calls, SQL injection and drive-by downloads [2]. Cloud networking will not change the fact that vulnerabilities will continue to exist, and that attackers will continue to exploit them. If anything, the concentration of massive amounts of computing power and data will make these targets more visible and more attractive.

This paper is organised as follows. Section 2 presents the concepts of cloud networking and cloud computing in more details. Section 3 explores the security issues when implementing the cloud networking vision. Section 4 presents closing remarks and summarises the next steps of this research project.

2 Cloud Networking

Cloud computing has gathered a lot of attention in recent years from parties across the computing and communication industries including vendors, network operators, and service providers. The service utility business model on which cloud computing is based is far from new. In 1961, Prof. John McCarthy was one of the first to introduce it by the claim that computer time-sharing technology might lead to a future in which computing power and even specific applications could be sold through the utility business model, i.e. water or electricity [3].

The existence of the internet and web technologies, and the introduction of infrastructure virtualisation has enabled the current realisation of that vision. Separation of the service provider from the infrastructure provider, is making it easier to generate new services on-line and to scale those services as demand dictates. For the service provider this reduces capital and operational expenditure, and financial risk, as they pay for access to resources on an as-needed basis, with little or no lead time to change capacity. For the infrastructure provider this gives the opportunity to build large infrastructures that benefit from economies of scale [4] and amortise the costs across the workload of multiple customers.

2.1 Virtualisation Technology Supporting Cloud Computing

Today's infrastructure-as-a-service (IaaS) is built on server virtualisation (virtual machine hypervisors such as Xen [5] or VMWare [6]), network virtualisation (implemented in network equipment or distributed routers such as [7]), and storage virtualisation (including network attached storage arrays or storage services such as Amazon's Elastic Block Store [8]). Data centre management systems deploy and manage virtual machines, networks, and data stores to construct any infrastructure topology required by the customer by dynamically re-configuring the virtualisation layers. These virtualisation techniques are now so common place that hardware support has been introduced to standard server chip sets by vendors such as Intel (VT-x [9]) and AMD (AMD-V [10]).

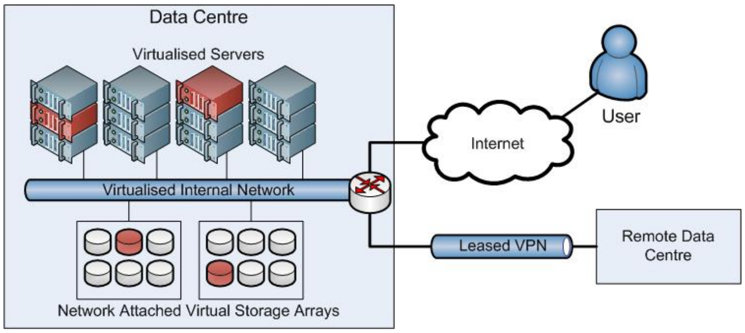

Fig. 1. Cloud Physical Infrastructure Architecture

The IaaS business model drives infrastructure providers towards a centralised architecture, as depicted in Figure 1. Very large data centres located near low cost power, land, and labour result in the lowest costs for the provider. However, the global nature of the business introduces opposing factors. From a regulatory perspective, the location of a data centre determines in part the legal jurisdiction that applies to hosted services (e.g. USA Patriot Act [11]). The use of the services can restrict their location or transfer of data (e.g. EU Data Protection Law [12]). From a technical perspective load, data transfer or disaster-tolerance may require multiple geographical locations. Processing load and data transfer is typically dealt with by parallel implementations of the service, each located geographically near the users. Disaster- tolerance requires replicating services in geographically diverse sites. As a result of these driving factors, today's cloud infrastructure providers operate a few, very large data centres, located in a select number of geographical locations.

Connectivity between data centres owned by a single provider is usually implemented by leased virtual networks providing guaranteed, but static quality of service for the IaaS owner. Connectivity between the data centre and the IaaS user is generally handled by the open internet. As such, the user's network experience is based on access to a shared medium, which is not under control of the cloud provider.

Although it is possible to scale the virtual infrastructure implemented by a IaaS provider, it is not possible to scale the connectivity to that infrastructure. Recently IaaS providers have added VPN tunnelling connectivity for their customers based on secure connection-oriented protocols such as IPsec (e.g. Amazon Virtual Private Cloud [13]). This allows, for example, the creation an IT infrastructure in the cloud that is connected to the site network of an enterprise in a way that enables them to use their own address space and network services across both. However, they are still subject to the limitations of bandwidth, jitter, and latency offered by their internet service provider and lack of support for dynamic provisioning.

The class of applications that are currently deployed in cloud infrastructures are those that are suited to this architecture, for example: batch processing, such as large scale simulations or graphics rendering, on-line web services, and hosted IT systems. Where sensitivity to network performance is an issue, such as content delivery [14], it is still necessary for the service provider to own the infrastructure or to enter into a long term contractual engagement with the infrastructure provider. The network components and topology of these services are still largely static.

2.2 Virtualisation Technology Supporting Cloud Networking

Network virtualisation brings a missing piece to the cloud computing puzzle. Virtual networks are not at all new in themselves; [15] provides a survey of technologies used at various layers. A number of network virtualisation architectures and frameworks have been proposed in the literature, including VINI [16], CABO [17], 4WARD VNet [18] and FEDERICA [19], to offer customised virtual networks with end-to-end control.

The possibility to specify and instantiate networks on demand and in useful time is one of the great advantages of network virtualisation. Virtual networks can be freshly created according to the different requirements, such as bandwidth, end-to-end delay, security, and protocols. Network virtualisation brings other advantages into stage. Such as the ability to reconfigure the network in real time without losing connectivity, to change the physical path or even to move one or more virtual nodes from one place to another [20].

Cloud networking extends network virtualisation beyond the data centre to bring two new aspects to cloud computing: the ability to connect the user to services in the cloud, and the ability to interconnect services that are geographically distributed across cloud infrastructures. These transform the cloud architecture of Figure 1 into that shown in Figure 2. Cloud networking users would be able to specify their needed virtual infrastructure and the desired networking properties to access these resources. Users would be able to specify how their infrastructure should be distributed in space and how it should be interconnected. They would be able to do this dynamically, on-demand, and through a single control interface.

The cloud paradigm has also encouraged the use of service automation. Applications running in cloud infrastructure can be programmed to monitor their own load and resource usage and dynamically scale themselves according to demand without the intervention of a human operator. Similarly, IaaS management systems optimise the use of physical resources by selectively deploying and migrating virtual machines. By introducing virtual networks to the same control plane the user and provider can make optimisation decisions based on network utilisation as well.

A new class of applications will be introduced to the world of cloud computing bringing with them new requirements. In some cases it may be more appropriate to deploy processing and storage functions across a network than to centralise

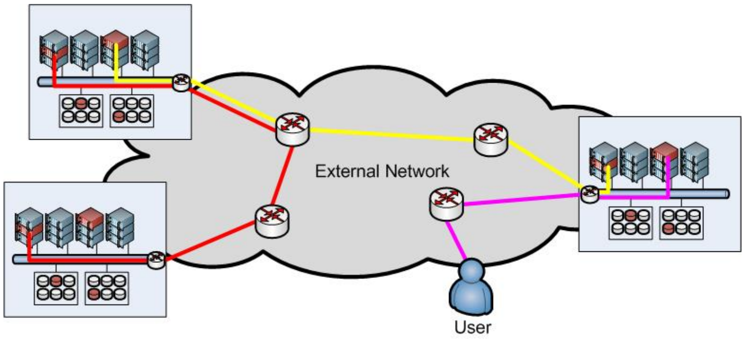

Fig. 2. Dynamic Virtual Networks Connecting a Distributed Service

processing and storage in a single location. The previously mentioned content distribution services might fit this approach for example.

We anticipate that a wider range of trade-offs between costs and performance requirements will lead to a wider range of deployment options. To enable these new possibilities, it is critical that we understand the security implications and build appropriate mechanisms into the technologies we develop.

Cloud networking puts at stage several security issues. In the following section we tackle them as a four dimensional problem.

3 Security Problem Space

It is anticipated that security is one of the major factors influencing the acceptance for cloud computing in practical application domains, especially when sensitive information shall be brought into the cloud or IT governance requires an elaborated control regarding the (legal) liability of computing in clouds [21]. From a users' perspective the security topics distinguish infrastructure security, platform and application security, the security of the management processes and finally compliance and governance [22].

The strength of the solutions that address these topics can be distinguished by the extent to which security objectives are met: who is allowed to do what (authentication & authorisation), how are system components and content protected (availability, confidentiality & integrity), how can the fulfilment of security properties be validated and checked (auditing), how can the cloud provider prevent others from doing forbidden things (misuse protection).

Cloud networking adds new security challenges to the cloud computing security issues, arising from additional networking capabilities. On the other hand, there are indications that cloud networking can potentially improve control over the cloud computing deployment model, thus solving the security challenges that impact acceptance of this technology. The following is a description of the security problem space as seen by the authors of the SAIL project at project start.

3.1 Information Security in Clouds

Information security relies on the classical three pillars, confidentiality (information should not be disclosed to unauthorised third parties), integrity (information should not be transformed without evidence of the transformation), and availability (information should not be withheld from rightful access).

The cloud scope adds a significant dimension in the mixing of code and data. Cloud users will need to ship code for execution on their data to cloud providers. Cloud providers will in turn ship code to users to easily manipulate the data. This is exemplified by the current rapid development of AJAX-based web services. Yet, this mixture of code and data is one of the major causes of malware infection, as it becomes extremely challenging to distinguish code from data and qualify the acceptability of both.

3.1.1 Trust in an Adversarial Environment Cloud environments are by their very nature adversarial. Cloud *providers* balance the needs of their multiple users, and attempt to monetise by-products of their activity. Cloud *users* strive to obtain the cheapest possible services, while requesting services of high quality, and respecting their privacy. *Attackers*, who have become very skilled at operating huge botnets (which can be seen as the first large scale clouds), will attempt to either access the information available in cloud, or avail themselves to this processing power free of charge. All actors thus have their own trust objectives, implemented in their security policies.

This adversarial setting promotes the use of security policy negotiation systems [23]. To maintain their trust relationship while ensuring sufficient flexibility to share resources, users and providers need to dynamically negotiate security policies, balancing between operational trade-offs such as cost and response time. This need will further develop as cloud providers aggregate and weave together complex service infrastructures federating many actors, creating the need for flexible and automated security policies aggregators and negotiators.

3.1.2 Confidentiality of Information and Processes One of the most effective ways to maintain integrity and confidentiality of information is encryption. While encryption in its current form is sufficient for data storage and transport, it fundamentally prevents data processing. Thus, sending encrypted data to cloud providers for processing is quite useless. This challenge has been met by homomorphic cryptography (HC). Homomorphic cryptography basically ensures that operations performed on an encrypted text basically still allows to retrieve by decryption the processed text in plain text. A solution under ideal circumstances has been presented recently [24, 25], and it is one of the very few HC systems that cryptographers have produced. However practical application is still far away since the computational effort required to retrieve the results of the computation is still too high and thus HC remains of theoretical value only for the coming years.

As a result, users will not have a cryptography solution that allows them to rely on information confidentiality and integrity when providing code and data to an arbitrary cloud. They have no means to ensure that their data is not misused. Until HC provides a formal solution to this issue, we need to rely on audit traces to assert "after the fact" usage control demonstrating that data and code have not been misused by service providers and cloud users. These audit traces can be part of a security policy specification, and can be supported for example by the OrBAC (Organisation-based access control) language. Further, we do not know yet if other solutions, like for example, watermarking will be portable to the cloud computing world and if their properties will be preserved in this world.

3.1.3 Policy Models and Policy Enforcement The currently available security policy models are not sufficiently flexible. For example, the OrBAC model [26], one of the recent attempts to further develop the classic RBAC model, introduces *organisations* and *contexts* in addition to the classic notion of roles. Both concepts are extremely useful to define security policies that span organisational boundaries (in our context multiple users sharing a cloud provider, or a federation of cloud providers uniting for a specific service) or security policies that are flexible according to environmental conditions (for example service load or cost). However, the combination of organisations and contexts with negotiation remains largely an unsolved problem. The complexity of these policies has not been resolved either. Even in simple environments such as network firewall filtering, users have difficulties understanding the impact of filtering rules when the number of rules is large, or when multiple firewalls are traversed. We expect that this complexity may become a barrier to the deployment of cloud computing if these policies cannot be simply explained and proven to all parties.

Security policies need to be enforced. Technology for this enforcement is reasonably well established using Policy Enforcement Points (PEP) controlling access to resources. Network firewalls, web application firewalls, identity management systems, file system access controls are well-know entities with clear properties. Policy Decision Points (PDP) are in charge of managing such PEPs and taking over for complex access control requests.

However, there is no such clear picture for the cloud computing world. First, it is not known if the policy enforcement technology can be ported into the could world and how. Second, it is unclear if cloud computing enabling technologies, such as virtualisation, will bring new PEPs. Once this setting is clearer, we will need to define techniques that will weave PEPs and PDPs into a cloud service definition, and tools for verifying that the resulting "secure service" definition meets the security objectives of all parties. It specifically requires new tools that will enable partial verification of security objectives, so that all parties (users and providers) can reliably verify that their security objectives are met, without knowing the security objectives of the other parties.

3.2 Virtualisation Environment Threats

Analysis of security threats in virtualisation environments provides some insight about the challenges raised by virtualisation of computing resources and networking. Figure3 highlights security issues that emerge when a type II hypervisor is used and identifies six security threats. They can be classified into software level (1, 2, 3, 4 and 6) and system level (5) concerns.

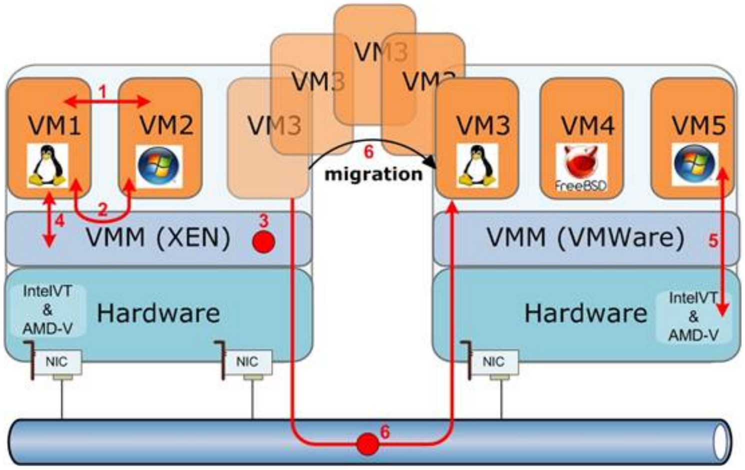

Fig. 3. Security Threats in Virtualised Environments

3.2.1 Isolation between Virtual Machines In this case, each virtual machine uses and only reads its allocated resources. For example, the memory management is subdivided into multiple levels (Hypervisor level, Host VM level and Guest VM level). The Hypervisor can read all the physical memory space. The Host Virtual machine (dom0 for XEN) can read all the memory except the memory allocated to the hypervisor. Guest Virtual Machines (domU for XEN) can only read their allocated memory. This isolation between different virtual machines is one of the main important roles of a hypervisor. As a solution, selection of the right hypervisor can ensure this isolation between virtual machines.

3.2.2 Information Theft through Malicious Use of Hypervisor To share physical resources, the hypervisor uses different techniques depending on the physical components to share. For example, to share physical network cards, the hypervisor (see the case of XEN at [27]) can use Bridged, NATed or Routed networking. In Figure 4, there are two bridges (*xenbr0* and *xenbr1*) that virtualize two physical network cards (*peth0* and *peth1*). The bridge *xenbr0* connects physical interface *peth0* to three virtual interfaces (*vif0.0*, *vif1.0* and *vif2.0*). Each virtual interface is connected to a virtual machine. In this configuration, despite the fact that all interfaces use the same bridge, it is necessary to ensure that a virtual machine cannot read the packets of the bridge that are sent to another virtual machine. This can be accomplished by the hypervisor or just by applying existing security solutions. To reduce the burden on the hypervisor in managing network I/O activities, manufacturers have since introduced Virtual Machine Device Queues (VMDq) [28] and Single Root Input Output Virtualisation (SR-IOV) [29]. Sorting data packets in the network silicon frees CPU cycles for application processing instead of network I/O processing. These new technologies introduce the additional requirement of securing, protecting and isolating also the network card virtualisation.

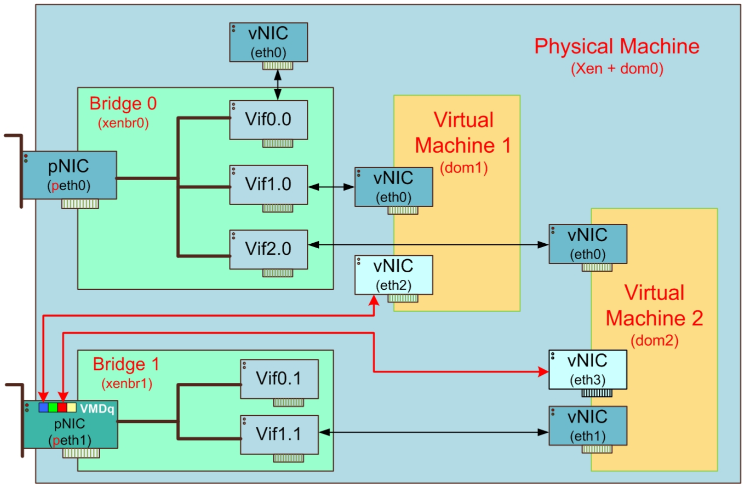

Fig. 4. Bridges Sharing Physical Network Cards

3.2.3 Untrusted Hypervisors If the owner of the physical machine wants to read and steal the data of virtual machines, she or he can do it using the hypervisor (an untrusted hypervisor). In this case, each user of a virtual machine needs to have a solid contract with the owner of the physical machine. Having a contract is a good and necessary thing, but it is imperative that virtual machines use their own mechanisms to secure themselves. For example, encrypting a virtual machine is a potential solution for this kind of problem.

3.2.4 Untrusted Virtual Machines It is always possible to have a contract to build some trust between the user of the virtual machine and the owner of the hypervisor. However, there is still the problem of the impossibility to have any idea about the other virtual machines than could be deployed in the same physical machine. A virtual machine can try to get control of the hypervisor using software related security holes without informing the hypervisor owner. Then, this virtual machine can get partial or total control of the physical machine. Technically, this is a similar situation to the untrusted hypervisor scenario. In this situation, it is possible to apply the same security solutions as in the untrusted hypervisor case.

3.2.5 Untrusted Virtual Machines Misusing Hardware Virtualisation Functionality To increase the performance of virtual machines in a virtualised environment, different functionalities (dedicated to virtualisation) have recently appeared in the architecture of physical components. As an example, new instruction sets (IntelVT-x [9, **?**] or AMD-V [10]) have been introduced in the most recent processors. With these functionalities, a virtual machine can send instructions directly to the processors by bypassing the hypervisor.

All the previously mentioned security problems can be solved using or adapting existing techniques. However, with these new types of security problems related to the hardware, a new philosophy and family of problems appear. Examples of systems that are sensitive or subject to these security threats are SubVirt and BluePill [31, 32].

3.2.6 Unsecure Network Transfer on Inter Device Migrations In a virtualised environment, a virtual machine can migrate from a physical machine to another. This migration through the network can use traditional or new protocols to attack the system. It is imperative to protect this migration by using or adapting existing techniques to prevent attacks on migration control mechanisms, transactions and protocols. This sixth identified threat is central in SAIL that focusses on cloud networking. This key aspect is tightly linked to auto-scaling and elasticity properties of clouds. In addition, there is a need for virtual firewalls for isolating dynamic VPNs and virtual networks allocated on the fly and on demand, to create dedicated flash slices.

3.3 Communication Security

Communication between virtual infrastructure as well as the distribution of virtual infrastructures generate traffic in the network, which has to be secured. The following Section 3.3.1 shows the challenges of securing the communication between virtual components, while Section 3.3.2 shows the security challenges of cloud networking, i.e. moving virtual components in space, and its management.

3.3.1 Secure Virtual Networking In addition to cloud computing, virtual networking introduces new security challenges by enabling communication between different virtual components. From a virtual network user's perspective the network might be private while in reality the communication itself occurs via a public infrastructure. Therefore, mechanisms to secure this communication (e.g., by encryption) have to be established. One option is to do it in each virtual component, which means that the virtual network customer has to care for securing the communication. Another option is to provide secured communication as a service by the virtual network provider, which means that the communication is secured by default and transparent to the customer.

Besides securing the communication itself in virtual networks, the management of the communication also has to be secured. By virtualising networks and network components new attacks arise and need to be handled. Due to the abstraction layer introduced by virtualisation, existing techniques might not be applicable or have to be adjusted or extended to fit this new setting. Especially the integrity of the virtual network topology and components, as well as the security of routing in these networks need to be addressed.

Additionally, similar challenges as in cloud computing also exist in virtual networks. This includes how the virtual network provider guarantees a certain network capacity to a customer, how the access to this virtual network is controlled and how the virtual network usage is accounted for.

3.3.2 Secure Management of Cloud Networking For the management of cloud networking access to the physical infrastructure and to the network properties is needed. This access should be implemented as a single interface, where a user can specify several parameters on-demand.

By the combined access to the physical virtualisation infrastructure and the network infrastructure new attacks arise. One challenge is to define rules for accessing the management interface and how to implement these rules. Also policies for moving virtual infrastructures in space need to be distributed. These policies might define to which location (legal space) a virtual infrastructure is allowed to move, as the location of the physical infrastructure determines the legal restrictions that apply to the virtual infrastructure (e.g., USA Patriot Act).

3.4 Misuse of Cloud Networking Capabilities

The ability of cloud computing and cloud networking to allocate computational resources on demand can also be misused, e.g. for DoS attacks, spamming and providing illegal content. Attacks that use cloud infrastructures are already known today. One example is Zeus "in-the-cloud" [33] where the command and control of a botnet was located at the Amazon EC2.

Auditing can help to detect these kind of attacks, e.g., by looking for fast fluxing or domain fluxing. The challenge of automated detection of attacks is to distinguish misuse from legitimate use. Trying to find anomalies might be one way to solve this problem. If a misuse can be detected the attack can simply be interrupted by discontinuing the virtual infrastructure, which is involved in the attack.

By introducing cloud networking no new threats are added to those already known from cloud computing. Therefore, countermeasures for misuse of cloud networking can be adapted from cloud computing.

4 Conclusion and Future Work

This position paper introduces the cloud networking specific security challenges that will be addressed in the SAIL project. These challenges can be grouped into protection of cloud content, secure virtualisation technology, distribution transparency control and secure operations. There are clear benefits that come with cloud networking for cloud users and operators. Also operators have prospect to support effectively cloud operators with their available network and transport capabilities for the benefit of end users. The road may even be open for further scenarios, e.g., connecting multiple clouds or introducing more heterogeneity, which in turn will increase the complexity in multilateral security. Both cloud computing and virtual networking have each their own security challenges, the ones presented here have to be considered for securing and protecting cloud networking that seeks technical solutions to ensure acceptance of this new concept.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the SAIL project.

References

1. Sail project website. URL `http://www.sailproject.org/`
2. Provos, N., Rajab, M.A., Mavrommatis, P.: Cybercrime 2.0: When the cloud turns dark. Queue **7**(2), 46–47 (2009). DOI http://doi.acm.org/10.1145/1515964.1517412
3. McCarthy, J.: MIT Centennial Speech of 1961 cited in Architects of the Information Society: Thirty-five Years of the Laboratory for Computer Science at MIT. SL Garfinkel Ed (1999)
4. Armbrust, M., Fox, A., Griffith, R., Joseph, A.D., Katz, R.H., Konwinski, A., Lee, G., Patterson, D.A., Rabkin, A., Stoica, I., Zaharia, M.: Above the clouds: A berkeley view of cloud computing. Tech. Rep. UCB/EECS-2009-28, EECS Department, University of California, Berkeley (2009)
5. Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I., Warfield, A.: Xen and the art of virtualization. In: SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles, pp. 164–177. ACM, New York, NY, USA (2003)
6. Vwmare. URL `http://www.vmware.com`
7. Edwards, A., Fischer, A., Lain, A.: Diverter: A new approach to networking within virtualized infrastructures. Tech. Rep. HPL-2009-231, HP Laboratories (2009)
8. Amazon elastic block store. URL `http://aws.amazon.com/ebs/`
9. Intel virtualization. URL `http://www.intel.com/technology/virtualization/`
10. Amd virtualization (amd-v) technology. URL `http://sites.amd.com/us/business/it-solutions/virtualization/Pages/amd-v.aspx`
11. Fraser, D.: The canadian response to the usa patriot act. Security Privacy, IEEE **5**(5), 66 –68 (2007)
12. Directive 2002/58/ec of the european parliament and of the council of 12 july 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (directive on privacy and electronic communications). In: Official Journal of the European Union, L201, pp. 0037–0047 (2002)
13. Amazon virtual private cloud. URL `http://aws.amazon.com/vpc/`
14. Pallis, G., Vakali, A.: Insight and perspectives for content delivery networks. Commun. ACM **49**(1), 101–106 (2006)
15. Chowdhury, N.M.K., Boutaba, R.: A survey of network virtualization. Computer Networks **54**(5), 862 – 876 (2010)
16. Bavier, A., Feamster, N., Huang, M., Peterson, L., Rexford, J.: In vini veritas: realistic and controlled network experimentation. In: SIGCOMM '06: Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 3–14. ACM, New York, NY, USA (2006)

17. Feamster, N., Gao, L., Rexford, J.: How to lease the internet in your spare time. SIGCOMM Comput. Commun. Rev. **37**(1), 61–64 (2007)
18. Schaffrath, G., Werle, C., Papadimitriou, P., Feldmann, A., Bless, R., Greenhalgh, A., Wundsam, A., Kind, M., Maennel, O., Mathy, L.: Network virtualization architecture: proposal and initial prototype. In: VISA '09: Proceedings of the 1st ACM workshop on Virtualized infrastructure systems and architectures, pp. 63–72. ACM, New York, NY, USA (2009)
19. Federica: Federated e-infrastructure dedicated to european researchers innovating in computing network architectures. URL `http://www.fp7-federica.eu/`
20. Wang, Y., Keller, E., Biskeborn, B., van der Merwe, J., Rexford, J.: Virtual routers on the move: live router migration as a network-management primitive. SIGCOMM Comput. Commun. Rev. **38**(4), 231–242 (2008). DOI http://doi.acm.org/10.1145/1402946.1402985
21. Brunette, G., Mogul, R.: Security guidance for critical areas of focus in cloud computing v2.1. Cloud Security Alliance (2009)
22. Streitberger W. Ruppel, A.: Cloud computing security - protection goals, taxonomy, market review. Tech. rep., Institute for Secure Information Technology SIT (2010)
23. Abi Haidar, D., Cuppens-Boulahia, N., Cuppens, F., Debar, H.: XeNA: an access negotiation framework using XACML. Annals of telecommunications **64**(1), 155–169 (2009)
24. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing, pp. 169–178. ACM, New York, NY, USA (2009). DOI http://doi.acm.org/10.1145/1536414.1536440
25. Gennaro, R., Gentry, C., Parno, B.: Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In: Proceedings of the International Cryptology Conference (CRYPTO) (2010)
26. Abou El Kalam, A., Baida, R.E., Balbiani, P., Benferhat, S., Cuppens, F., Deswarte, Y., Miège, A., Saurel, C., Trouessin, G.: Organization Based Access Control. In: 4th IEEE International Workshop on Policies for Distributed Systems and Networks (Policy'03) (2003)
27. Xen networking blog (2010). `http://wiki.xensource.com/xenwiki/XenNetworking`
28. Shefali Chinni and Radhakrishna Hiremane: Virtual machine device queues (vmdq) - white paper (2010). `http://software.intel.com/file/1919`
29. Pci-sig single root i/o virtualization (sr-iov) support in intel virtualization technology for connectivity - white paper (2008). `www.intel.com/network/connectivity/solutions/SR-IOV-046NTL_Whitepaper_061308.pdf`
30. Uhlig, R., Neiger, G., Rodgers, D., Santoni, A., Martins, F., Anderson, A., Bennett, S., Kagi, A., Leung, F., Smith, L.: Intel virtualization technology. Computer **38**(5), 48–56 (2005)
31. Price, M., Partners, A.: The Paradox of Security in Virtual Environments. Computer **41**(11), 22–28 (2008)
32. King, S.T., Chen, P.M., Wang, Y.M., Verbowski, C., Wang, H.J., Lorch, J.R.: Subvirt: Implementing malware with virtual machines. Security and Privacy, IEEE Symposium on **0**, 314–327 (2006). DOI http://doi.ieeecomputersociety.org/10.1109/SP.2006.38
33. CA Community Blog: Zeus "in-the-cloud" (2009). `http://community.ca.com/blogs/securityadvisor/archive/2009/12/09/zeus-in-the-cloud.aspx`