

Dynamic Topology Extraction in Cloud Infrastructures

Pernelle Mensah

► **To cite this version:**

Pernelle Mensah. Dynamic Topology Extraction in Cloud Infrastructures. Second workshop on Security in Clouds (SEC2), Jul 2016, Lorient, France. hal-01399251

HAL Id: hal-01399251

<https://hal.inria.fr/hal-01399251>

Submitted on 18 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Dynamic Topology Extraction in Cloud Infrastructures

Pernelle Mensah
pernelle.mensah@nokia.com

*Nokia / CentraleSupélec / Inria
Secured Cloud Networking (SCN)
7 Route de Villejust, 91620 Nozay (France)*

Abstract

To determine the threat exposure of a virtualized environment, attack graphs generation, coupled to a risk-based assessment can be used. The first roadblock to lift to that end is the extraction of the topology. We will present in this paper the strategy we intend to use to obtain a near real-time view of the connectivity existing in a virtual infrastructure.

1 Introduction

Communication infrastructures are regularly targeted by attackers willing to gain access to the value they contain. Viable and efficient security supervision mechanisms are hence of the utmost importance for companies in this digital era, in order to trigger a swift remediation in case of intrusion detection. To address this issue, approaches such as attack graphs generation, coupled to a risk-based assessment have been used to provide an insight into a system's threat exposure [10, 12, 14]. After identification and evaluation of existing attack paths allowing a compromise of the system, measures can be deployed for mitigation. However, as corrective actions are often constrained by operational requirements that can delay their enforcement, (e.g. need to ensure a continuity of service, cost of measures to deploy...), a balance has to be found with the security policy enforced.

Such approaches have yet to be applied to a virtualized context, as these environments, besides sparking a growing interest, also give rise to new challenges. On one hand, the very elasticity and dynamism that draw to the Cloud leads to more volatile customers' infrastructures and more mobile virtual components. Furthermore, new attack scenarios have to be considered since logical boundaries between a hypervisor and its virtual machines (VMs) can be circumvented by a malicious VM to access the host system or other co-located instances

[2, 1]. On the other hand, virtualized environments abide by security policies that originate from at least two governances: the one of the provider, and the one of the tenants (customers sharing the same hardware in a cloud environment). Taking these considerations into account, several steps are involved in the realization of a dynamic risk management solution for the Cloud:

- Extracting the topology: knowledge gained concerning connectivity and reachability within the tenants' infrastructure and with the provider system is a basis for an accurate view of the assets and their interconnectedness;
- Translating that topology into an attack graph allows for the identification of feasible attack paths in the infrastructure;
- Determining the risk, based on the attack graph obtained, to guide the administrators in the decision process;
- Activating the proper response (based on previous established order) in adequation not only with the security policy of the provider, but also the one of the related tenant and potentially its co-located neighbours.

In the remainder of this article, we elaborate more on the first item of the solution: the extraction of the topology. The goal is to present the strategy we intend to use to obtain a near real-time view of the connectivity existing in the virtual infrastructure.

2 Motivating Use Case

2.1 Community Cloud of Telecommunication Providers

The use case considered stems from our research environment, i.e. a company providing services es-

essentially to telecommunication providers, and combining cloud, SDN (Software-Defined Networking) and NFV (Network Functions Virtualization) solutions. NFV enables a "softwarization" of the network: networking functions previously bound to purposed-built hardware (firewalls, routers, load-balancers, encryption...) can now be implemented as software and supported by off-the-shelf equipments. Additionally, combined with cloud computing and SDN for dynamic network configuration, these functions can be deployed into virtual machines, allowing the telecommunication providers to benefit from scalability, better load management, economies of scale and faster innovation. Because, to some extent, telecommunication providers depend on analogous functionalities to provide their services, and share similar concerns regarding security and compliance, they can be aggregated into a community cloud [9]. We thus design a use case based on a company operating a community cloud, and hosting virtualized network functions for its tenants, i.e telecommunication providers. The




Figure 1: Example Architecture

cloud provider's infrastructure consists of a cloud management system, coupled with a SDN controller and Virtual Machine Introspection (VMI) mechanisms. Figure 1 shows an example of the architecture, with tenants leveraging the physical infrastructure of the cloud provider to supply their services.

2.2 A Hierarchy of Dynamics

Given the context previously introduced, we can assume that in such clouds, a rough uniformity can be deduced regarding the characteristics of the tenants' infrastructures and a first assumption can

be made regarding the different levels of dynamicity in the considered infrastructure. Indeed, if we consider the tenants to be telecommunication providers, their architectures can be considered relatively stable from the standpoint of the network connectivity. Being managed by telecommunication companies to provide services to their customers, major network changes are first tested for a proper integration with production environment to avoid loss of service. On the other hand, the transfer of VMs from one physical machine to another, their addition, deletion or extinction occur at a much faster rate, according to the cloud provider's and tenants' operational needs, as well as potential SLAs negotiated between the parties. Indeed, when facing peak demands, telecom providers may need to spawn new instances to host their network functions, instances that would be terminated as the traffic drops down. We hence hypothesize a hierarchy of dynamics, the rate of changes in network configuration being lower than the one in the location and state of the VMs. These are properties that will be considered for an optimisation of the topology generation, by an identification of operations resulting in a complete re-computation of the topology and those requiring only a slight update.

2.3 Challenges to Topology Extraction

Besides the dynamic nature of the Cloud, lack of knowledge regarding tenants' custom configurations can also hinder the topology acquisition. Indeed, the use of a cloud management system make tenants' infrastructure details readily available, exactly as they are provisioned by the tenants on the provider's platform; yet, no hypothesis can be made upon connectivity restrictions implemented by the tenants within their infrastructures, e.g. custom firewall instantiation within a VM. Furthermore, the combination of Cloud and SDN poses an additional challenge, since an essential property of SDN is enabling the administrators to deploy applications on the SDN controller. Those applications can then, according to the programmed logic, reactively modify the flow rules on the switches and directly influence the topology. They can add or remove links, without providing any feedback to the Cloud Management System, leading to an inconsistent state of the topology processed via that source. We thus need to closely monitor changes introduced by SDN applications on the topology.

3 An Approach for Topology Extraction in the Cloud

3.1 Background

The topology can be retrieved by following either a static or a dynamic approach. Madi and al. [8] extracted data from a cloud management system, a SDN controller and virtual switches in order to check (among other things) whether the topology view in the cloud management system matches the actual implementation. However, they performed their analysis on a snapshot of the infrastructure and did not address dynamic changes. Probst [13] investigated the correct placement of security mechanisms in virtual infrastructures, by auditing a clone of the targeted virtual environment. Copying the initial network configuration and virtual firewalls, and replacing the virtual machines by lighter, custom versions allowed to rebuild the customer’s infrastructure, while maintaining the confidentiality of her data. In both research, the authors considered a frozen picture of the virtual infrastructure, in which a change called for a re-acquisition and re-processing of the data. To detect isolation failure in virtualized environments, Bleikertz and al. [6] performed a topology discovery phase, and a translation of the retrieved data into a graph. By using hypervisors’ APIs, they aggregated knowledge regarding the configuration of the virtual environment: physical machine, virtual machine, storage and network details. Rather than periodically refreshing the entire configuration details, they addressed the dynamic nature of that environment [7], by monitoring change events in the infrastructure to update specific portions of the graph. We also intend to monitor change events, but we go a little bit further by hypothesizing that some changes occur more frequently than others. That consideration can help the design of an efficient data processing strategy.

3.2 Defining Connectivity in a Virtualized Environment

Traditionally, the notion of connectivity in a telecommunication context is associated with the ability of a particular machine (either an endpoint or intermediary equipment) to reach another through the network. This is often referred to as *network connectivity*. However, as evoked in the introduction, cloud computing allows for the replacement of physical boundaries by logical ones, which can be by-passed, thus exposing the resident machines to intruders. Indeed, attacks can be performed to initiate VM-to-hypervisor or VM-to-VM

compromise [11], resulting in possible information extraction from the targeted tenant, or full control over the virtualization layer. As a consequence, aside from the network connectivity, a *local connectivity* also exists, representing the relationship between hypervisors and their hosted VMs. We can identify hypervisors and virtual machines as two distinct types of nodes. The expression of the various components connectedness can hence be represented as follows:

- A network connectivity between hypervisors;
- A network connectivity between virtual machines instances;
- A local connectivity between virtual machines’instances;
- A local connectivity between the hypervisor and its virtual machines’ instances.

3.3 Preliminary Experimentation

In this preliminary experimentation, our objective was to retrieve the local and network dependencies. A simple architecture of three servers was used for tests. Built upon KVM [3] as the hypervisor technology, Openstack [5] as the Cloud Management System and Opendaylight [4] as the SDN controller, one server was dedicated to Openstack control functions, the SDN controller and the networking service, while the other two were used for VM hosting. Openstack and Opendaylight were chosen due to their widespread use in the industry and the continuous efforts made to integrate these platforms. We considered the representation of information extracted from Neutron and Nova databases, respectively Openstack’s networking and compute services. Together they store knowledge about the virtual networks, the virtual machines and hypervisors involved, as well as the security rules applied to the guest machines (Openstack’s firewall rules). Objects in the local connectivity are hypervisors and VMs. On the other hand, the network connectivity is comprised of VMs, virtual networks, virtual routers and virtual security groups as objects. As an initialization phase, for each hypervisor in Openstack database, we associate its running VMs (local connectivity). Each VM is linked to its security groups and directly connected networks, each network potentially having an entry for routers connected to it. That initialization allows to have a first snapshot of the state of the infrastructure. Each change of VM or network state creates a notification that is queued in Openstack messaging service for processing; in

order to dynamically update the state of our objects, we intercept and process these events. In our first approach we registered the resulting topology (local and network) into a single formatted file that had to be parsed and updated upon any given event. However, as described in the use case, due to our context, the virtual networks have a slower evolution than the local connectivity. Events such as VM migration or instantiation do not affect the structure of the network designed by the tenants. Considering the addition of physical nodes by the providers, the same can be assumed for the physical network, at least the segments connecting the nodes. Due to this finding, we splitted the file in two, one for each type of connectivity, and reduced the computation overhead when a modification occurs depending on its impacts: either on local or on network connectivity.

4 Conclusion

In this article, we presented our early thoughts regarding the extraction of the topology in a virtualized environment. An experimental context has been introduced. Dynamic aspects, which represent a strong suit of the Cloud, but also the main source of concern for an acquisition of a complete view of the tenants' infrastructures, have been categorized according to change rates. This represents the first step in the elaboration of a general solution for dynamic risk management applied to modern communication infrastructures. From that preliminary study emerged a physical and a virtual strata, which we intend to leverage for an optimized topology extraction and afterwards a multi-layer attack graphs generation. The ultimate intent being to guide cloud providers and tenants in their security decision process, by prioritizing attack paths discovered in the system and applying the proper security policy.

References

- [1] Cve-2015-8555. <https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2015-8555>. Accessed 23rd May 2016.
- [2] Cve-2016-3710. <https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2016-3710>. Accessed 23rd May 2016.
- [3] Kvm. http://www.linux-kvm.org/page/Main_Page. Accessed 27th May 2016.
- [4] The opendaylight platform. <https://www.opendaylight.org/>. Accessed 27th May 2016.
- [5] Openstack open source cloud computing software. <https://www.openstack.org/>. Accessed 27th May 2016.
- [6] Sören Bleikertz, Thomas Groß, Matthias Schunter, and Konrad Eriksson. Automated Information Flow Analysis of Virtualized Infrastructures. In *Proceedings of the 16th European Conference on Research in Computer Security, ESORICS'11*, pages 392–415, Berlin, Heidelberg, 2011. Springer-Verlag.
- [7] Sören Bleikertz, Carsten Vogel, and Thomas Gross. Cloud Radar: Near Real-Time Detection of Security Failures in Dynamic Virtualized Infrastructures. In *Annual Computer Security Applications Conference (ACSAC)*, 2014.
- [8] Taous Madi, Suryadipta Majumdar, Yushun Wang, Makan Pourzandi, and Lingyu Wang. Auditing security Compliance of the Virtualized Infrastructure in the Cloud: Application to OpenStack. In *6th ACM Conference on Data and Application Security and Privacy ACM CODASPY 2016*, 2016.
- [9] P. Mell and T. Grance. The NIST Definition of Cloud Computing, September 2011.
- [10] S. Noel, S. Jajodia, and A. Singhal. Measuring security risk of networks using attack graphs. *International Journal of Next-Generation Computing*, 1(1), 2010.
- [11] J. Oberheide, E. Cooke, and F. Jahanian. Empirical exploitation of live virtual machine migration. In *BlackHat DC convention*, 2008.
- [12] X. Ou and A. Singhal. *Quantitative Security Risk Assessment of Enterprise Networks*. Springer, 2012.
- [13] T. Probst. *Evaluation et Analyse des Mécanismes de Sécurité des Réseaux dans les Infrastructures Virtuelles de Cloud Computing*. PhD thesis, Université Fédérale Toulouse Midi-Pyrénées, 2015.
- [14] L. Samarji, N. Cuppens-Boulahia, F. Cuppens, W. Kanoun, S. Papillon, and S. Dubus. Liccac: Assessing the likelihood of individual, coordinated, and concurrent attack scenarios. In *10th International Conference on Security and Privacy in Communication Networks (SecureComm 2014)*, sep 2014.