Security Architecture for Cloud Networking

Volker Fusenig and Ayush Sharma
Fraunhofer Research Institution
for Applied & Integrated Security
Parkring 4, 85748 Garching, Germany
Email: {firstname.lastname}@aisec.fraunhofer.de

Abstract—**Cloud computing offers reduced capital expenditure, operational risks, complexity and maintenance, and increased scalability while providing services at different abstraction levels, namely Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS). A new approach called cloud networking adds networking functionalities to cloud computing and enables dynamic and flexible placement of virtual resources crossing provider borders. This allows various kinds of optimization, e.g., reducing latency or network load. However, this approach introduces new security challenges. This paper presents a security architecture that enables a user of cloud networking to define security requirements and enforce them in the cloud networking infrastructure.**

I. Introduction

In recent times cloud computing has become more and more popular and is applied for various purposes. Cloud computing itself is in principle an abstraction of the physical infrastructure which is offered as cloud services to service users. The abstraction levels of these cloud services are Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS). Popular examples are GoogleDocs [1] for SaaS, Google's AppEngine [2] for PaaS, and Amazon's EC2 [3] for IaaS.

The cloud computing infrastructure is hosted in data centers. If a service user orders a cloud service, e.g., a virtual machine in Amazon's EC2, this virtual resource is placed on a physical infrastructure within a data center of the cloud operator. The virtual resource might be moved within the data center from one physical machine to another, e.g., due to maintenance reasons. Migrating virtual resources to physical machines located in other data centers of the same operator, or to physical machines of other operators automatically is not possible. However, there are some use cases where a flexible placement of virtual resources is needed, e.g., for optimization reasons (reducing costs or latencies when accessing the virtual resource).

The European project SAIL [4] investigates the combined management of cloud computing infrastructures (of different operators) and network infrastructures. This combined management enables the infrastructure service user to optimize various parameters, like costs for the virtual resource, latencies when accessing them. Additionally, this combined management enables cloud and network operators to enhance the work load of their infrastructure, e.g., by adapting pricing or by intelligent placement of virtual resources to reduce network load.

Unfortunately, this flexible placement of virtual resources introduces new challenges regarding security [5]. In the cloud computing world the service user checks the security level of a cloud operator manually if security relevant information is published by the operator. Only if the security policies of the service user are followed the service user moves his virtual resource to the cloud operator's infrastructure and it stays there until the costumer removes it manually.

In cloud networking, virtual resources are moved automatically from one operator's cloud infrastructure to another. Therefore, security checks must also be carried out automatically, in order to assure that an operator's infrastructure follows the service user's demands. In this paper we show an approach for automated security checks. In our approach a service user can define security requirements and a virtual infrastructure provider can describe its security functionality. A service provider moderates the placement of virtual resources, maps the security demands to security functionality, and if needed moves the virtual resources to another virtual infrastructure provider.

The paper is organized as follows. First, we introduce roles and use cases in cloud networking which will be used during the rest of this paper (see Section II). In Section III we show security challenges we want to solve with our approach and show how to define security parameters that are build up from security policies of service users and security mechanisms of virtual infrastructure providers. Based on this we introduce our approach for requesting security functionality based on service users' requirements. In the following Section IV we show how the security functions are added to the cloud network architecture and define functions that are needed to interact between architectural components. Afterwards, in Section V we give an overview on related work that is relevant for securing cloud networking. The last Section VI concludes the work and shows further working directions.

II. Roles and Use Cases

In this section we introduce the terminology for cloud networking as used in this paper. This consists of roles and actors in this environment. Furthermore, we show two example use cases for cloud networking to illustrate these roles and show the basic functionality of cloud networking.

978-1-4673-0009-4/12/$26.00 ©2012 IEEE

A. Roles

We will use the following terminology during the rest of this paper.

Service User: A service user is a person who is authorized to use the services provided by a service provider, e.g., a private user or an employee of an enterprise. He requests services by the service provider by using a request mechanism.

Service Provider: A service provider provides services to a service user. He is responsible to map the requested services of the service user to a virtual infrastructure provider. Especially, he has to take care of following the security demands of the service user. A service provider can also be at the same time a virtual infrastructure provider.

Virtual Infrastructure Provider: A virtual infrastructure provider provides virtual IT infrastructure, where services can be installed on. A service provider uses this virtual infrastructure for implementing and offering his services. Generally, the virtual infrastructure provider owns the hardware, where the virtual infrastructure is running on.

Virtual Resource: A virtual resource is a virtual processing, storing, or networking entity that is placed on a physical resource of a virtual infrastructure provider.

B. Use Cases

Travelling business man: In this use case we consider a business man (service user) who has only a lightweight portable device and uses this device to access his virtual machine (service offered by the service provider) which runs in the cloud network. The service user has several security demands on the virtual resource, in order to follow his company's security policies, e.g., that the virtual resource must be located in Europe or Australia and that the access to the data center where the virtual resource resides must be ISO 27001:500 [6] certified. His normal working place is in Europe. In order to reduce latency the service provider demands virtual resource in Europe by a virtual infrastructure provider which has physical resources in Europe that are ISO 27001:500 certified.

When the business man is on a travel in Australia the service provider takes care of moving the business man's virtual resource to Australia in order to keep latency down. Again, the service provider has to take care of choosing a virtual infrastructure provider that has ISO 27001:500 certified physical resources.

On another time the business man travels to the USA. In this case the service provider cannot move the virtual machine to the USA because of the security demands of the business man.

Cheap Processing and Storage In this use case we consider a service user that is interested in cheap processing and storing resources, e.g., a small company that needs from time to time some intensive calculations (e.g., rendering of videos). The service user demands processing power with the constraint that the service is operated in a data center that is ISO 27001:500 certified at the lowest price. In this use case the latency is not really important. For that reason the service provider takes the cheapest virtual infrastructure provider that is ISO 27001:500 certified. If the processing of the task takes longer the service provider may move the task to another virtual infrastructure provider if it becomes the cheapest one.

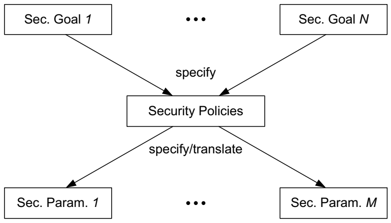

Fig. 1. Security parameters of service users

Reasons for diversities in prices of a virtual infrastructure provider might be the current work load (e.g., because of different time zones of service users and virtual infrastructure provider) or diversity in prices for energy (smart grid).

III. Security Challenge and Approach

One challenge in cloud computing is that the service user has some security requirements on the cloud infrastructure which he wants to use, e.g., in the first use case of Section II-B these security requirements base on security policies of the business man's company. Today, this is done manually in the way that the service user checks some security parameters of a virtual infrastructure provider and compares them with his security requirements. After a positive evaluation the service user moves his data or processes to the virtual infrastructure provider's place. In the case that security parameters of the virtual infrastructure provider change the service user has to check the parameters and security requirements again manually. The same holds if the service user wants to use the cheapest virtual infrastructure provider (see second use case in Section II-B). In this case the service user has to compare prices of different virtual infrastructure providers manually, check the security level of the cheaper virtual infrastructure provider manually before moving the resources, and move the virtual resources manually to the new place.

The cloud networking approach helps to distribute the virtual resources flexibly to different virtual infrastructure providers. The service provider takes care of optimization and harmonization of different parameters, e.g., latencies, cost, and network load. In the same way the service provider has to take care of respecting the security requirements of the service users. In the following two subsections, we first show how a service user defines security goals and how these goals are translated into security parameters. On the other hand, we show how a virtual infrastructure provider describes his security functionalities as security parameters. Second, we show the process of implementing the security requirements of a service user into the infrastructure of the virtual infrastructure provider by defining constraints on the resources of a virtual infrastructure provider.

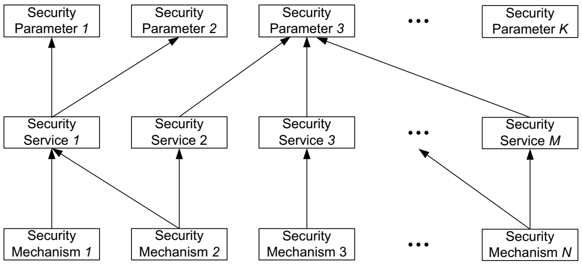

Fig. 2. Security parameters of virtual infrastructure providers

A. Extraction of Security Parameters

Figure 1 shows the process of extracting security parameters out of security goals of a service user. In a first step the service user defines security goals for his virtual resources, e.g., confidentiality of stored data and integrity of stored data. The service user expresses these security goals as a security policy, which is more or less a list of security requirements, e.g., containing the statement "data must be stored encrypted". This security policy is then translated into a list of security parameters. This translation step has two reasons: First, we can define a unique description language for security parameters in order to have means for comparing security requirements of a service user and security mechanisms of virtual infrastructure providers (see next section). Second, we can limit the number of potential security parameters to a list of predefined ones. E.g., a security policy saying "data must be AES encrypted" and a second policy saying "data must be encrypted" can both be described with a security parameter "encryption". This parameter might have the entries "encryption scheme" and "key lenght". In the first case the policy might results in "encryption scheme == AES, key length == all" and in the second case in "encryption scheme == all, key length == all". The definition of a description language is not part of this paper. Most likely the description language for security parameters will be based on XML (e.g., using VXDL [7]).

Using the same security parameters we want to describe the security functionality of a virtual infrastructure provider. Figure 2 shows how the security parameters are extracted from the security mechanisms installed at the virtual infrastructure provider's side. A security service (confidentiality service) is in this case an abstraction of one or more security mechanisms (e.g., AES encryption with 256 bit key length). As can be seen in the figure there might be security parameters which do not base on a security mechanisms. This can be the case where no technical mechanism is needed to implement this parameter, e.g., location of the virtual infrastructure provider's side.

As result we now have security parameters on service user's side describing his security requirement. On the other side we have security parameters describing the security functionality of a virtual infrastructure provider.

B. Approach

In the previous Section III-A we have seen how the security parameters are extracted from the service user and from the virtual infrastructure provider. The virtual infrastructure provider needs to commit his list of security parameters describing his security functionality to the service provider. The service provider stores for each virtual infrastructure provider he has contact to such a list of security parameters. Besides this, he also stores additional information on functionalities and available resources which is not part of this paper.

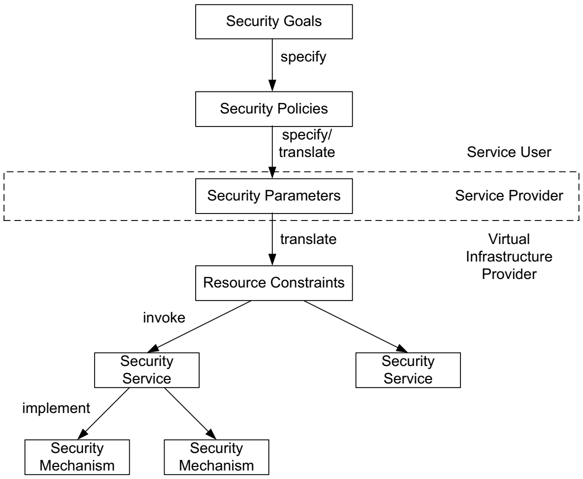

Fig. 3. Security approach

When a service user wants to have a virtual resource, e.g., cheap storage, he sends a request to the service provider (see second use case in Section II-B). Together with the request for virtual resources he sends a list of security parameters which needs to be followed. In the next step the service provider compares these security parameters with the security parameters of virtual infrastructure providers. If a virtual infrastructure provider has at least the security functionality as requested through the security parameters by the service user the service provider might invoke virtual resources at this virtual infrastructure provider. The decision on where to place virtual resources also depends on other parameters, e.g., price, which is not part of this paper.

After the service provider has chosen a virtual infrastructure provider that fulfils the security requirements (described as security parameters) of the service user he translates the security parameters to resource constraints for the virtual infrastructure provider. This step is needed because the service user does not need all security functionalities of the virtual infrastructure provider (e.g., only AES encryption with 256 bit key length and no DES encryption). Based on these resource constraints the virtual infrastructure provider invokes security services and mechanisms (AES encryption with 256 bit key length) for virtual resources.

IV. Architecture and Functions

The architecture and functions between the architectural elements are shown in Figure 4.

This architecture consists of three functional entities: the service user, the service provider, and the virtual infrastructure

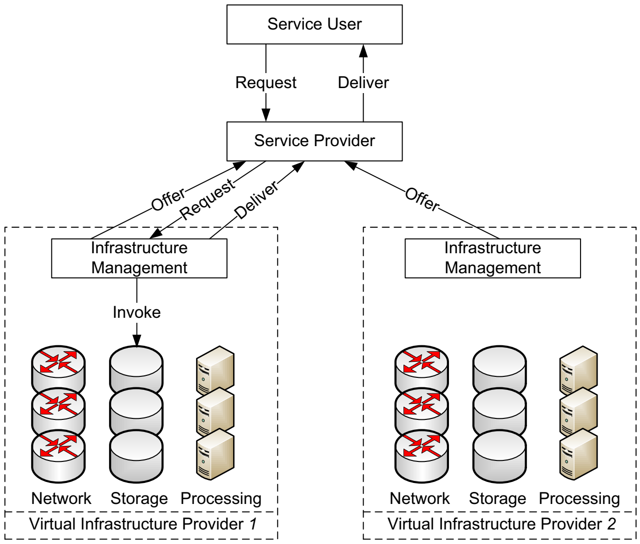

Fig. 4. Architecture and security function

provider. The figure shows the different functions which are used to interact between these entities. We classify the functions by the caller:

A. Service User Functions

Request Virtual Resources In this simplified architecture the only function a service user can call is a request for virtual resources from the service provider. This request contains overall goals, e.g., type of resource, amount, and optimization parameters (e.g., latency demands and price), and the security goals. The security goals are transmitted in form of security parameters as described in Section III-A. The translation of security goals to security parameters will be performed at the service user side. The translation step can be assisted by translation tools.

B. Service Provider Functions

Request Virtual Resources The service provider has also a function for requesting virtual resources. This function requests virtual resources from a virtual infrastructure provider. However, before sending the request to a virtual infrastructure provider the service provider maps the security parameters of the service user to the security parameters of a virtual infrastructure provider. Only if the virtual infrastructure provider provides the security functionality to fulfil the security demands of the service user the service provider is allowed to request the resources here. The request contains the same overall goals as the request function from the service user and additional security constraints which base on the security parameters of the service user.

Deliver Virtual Resources After receiving the access to the virtual resources from the virutal infrastructure provider the service provider forwards the access to the service user. Details on access control is not part of this paper and will follow in future work.

C. Virtual Infrastructure Provider Functions

Offer Virtual Resources and Security Functionality The virtual infrastructure offers his virtual resources and security functionality to the service provider. How the virtual resources are offered is not part of this paper. The security functionalities are described as security parameters (see Section III-A).

Invoke Virtual Resources and Security Functionalities When the virtual infrastructure provider receives a request for virtual resources from a service provider he invokes the virtual resources and activates the requested security functionalities (see request function of service procvider).

Deliver Virtual Resources After having invoked the virtual resources and security functionalities the virtual infrastructure provider sends the access to the virtual resources to the service provider.

D. Utilization

We show the utilization of the security architecture and its functions by applying it to the second use case of Section II where a service user is interested in cheap processing and storing resources.

In a first step the virtual infrastructure providers VP_1 and VP_2 **offer** virtual resources and security functionality to a service provider (see Figure 4). VP_1 offers storage resources and provides AES encryption with maximum key length of 256 bit as security functionality (security parameter). VP_2 also offers storage resources and provides DES encryption with 56 bit key length as security functionality.

The service user SU sends a **request** for virtual resources to the service provider. SU asks for 100GB of storage that is AES-encrypted with at least 128 bit key length (security parameter).

The service provider now detects that only VP_1 offers the adequate encryption for the request of SU. For that reason the service provider sends a **request** to VP_1 for 100GB of storage that is AES encrypted with 128 bit key length (resource constraint).

VP_1 **invokes** the virtual resource by allocating 100GB of storage, encrypting the storage with AES and 128 bit key length, and **delivering** the access to the virtual resource to the service provider. The service providers forwards the access to SU.

V. Related Work

As cloud networking is a new research topic performed in the SAIL project [4] there is no specific work on security in this field except a collection of security challenges in cloud networking [5] which is the basis of this work. However, the work presented in this paper is also related to traditional cloud computing. An overview on cloud computing security goals can be found in [8] and typical attacks on cloud infrastructure in [9].

There are several approaches on how to enforce security policies in cloud computing, e.g., secure selective sharing of resources based on new encryption schemes [10]. Basescu et. al. [11] propose a security management framework for cloud computing for defining and enforcing flexible security policies. However, this framework is only for a single provider and is not applicable to a multi-provider approach with the flexible distribution of virtual resources as needed for cloud networking. A common ground of their work and the work described in this paper is the definition and expression of security policies.

Iskander et. al. [12] propose a mechanisms for enforcing authentication policies in cloud computing. The CloudAudit group [13] is working on automation of audit, assertion, assessment, and assurance in cloud environments. We plan to add auditing, policy enforcement, and policy verification techniques to our architecture and will include results from these approaches in our work.

The basis for virtualization as used in cloud infrastructures are for example XEN [14] or VMWare [15]. Additionally, as cloud networking also needs the access to networking resources network virtualization techniques [16] are needed. There exist also hardware extensions for better performance in virtualized environments (e.g., for AMD Virtualization Technology [17] and Intel Virtualization Technology [18]). Virtualisation techniques and hardware acceleration techniques become important when implementing the security architecture.

VI. Conlusion

In this paper we presented a security architecture for cloud networking. This architecture helps in preserving the security goals of service users while at the same time benefiting from the flexible and dynamic placement of virtual resources at different virtual infrastructure providers. Key concepts of this architecture are the definition of unique security parameters for expressing security requirements and security functionality, the translation of security parameters in security constraints, and the management of service users and virtual infrastructure providers by service providers.

As further steps we plan to include the security functionality in the SAIL prototype. We plan to extend the architecture by auditing techniques so that a service user and a service provider are able to verify that a security constraint is followed by a virtual infrastructure provider. Furthermore, we plan to establish access control mechanisms for accessing virtual resources and making the flexible nature of the cloud networking infrastructure transparent for the service user.

VII. Acknowledgements

The authors would like to express their gratitude to the European Commission for its funding through the "Scalable and Adaptive Internet Solutions", SAIL Project (FP7-ICT-2009-5-257448).

References

[1] "Google Docs," July 2011. [Online]. Available: http://docs.google.com

[2] "Google App Engine," July 2011. [Online]. Available: http://code.google.com/appengine/

[3] "Amazon Virtual Private Cloud," July 2011. [Online]. Available: http://aws.amazon.com/ec2/

[4] "SAIL project website," July 2011. [Online]. Available: http://www.sail-project.eu/

[5] P. Schoo, V. Fusenig, V. Souza, M. Melo, P. Murray, H. Debar, H. Medhioub, and D. Zeghlache, "Challenges for cloud networking security," in *Mobile Networks and Management*, ser. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering. Springer Berlin Heidelberg, 2010.

[6] "ISO/IEC 27001:500 - Information security management systems - Requirements," July 2011. [Online]. Available: http://www.iso.org/iso/catalogue_detail?csnumber=42103

[7] G. P. Koslovski and P. V.-B. Primet, "Vxdl: Virtual resources and interconnection networks description language," *Engineering*, pp. 138–154, 2009. [Online]. Available: http://www.springerlink.com/index/N131016226414516.pdf

[8] A. Streitberger, W. Ruppel, "Cloud computing security - protection goals, taxonomy, market review," Institute for Secure Information Technology SIT, Tech. Rep., 2010.

[9] N. Provos, M. A. Rajab, and P. Mavrommatis, "Cybercrime 2.0: When the cloud turns dark," *Queue*, vol. 7, no. 2, pp. 46–47, 2009.

[10] S. D. C. d. Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, G. Pelosi, and P. Samarati, "Encryption-based policy enforcement for cloud storage," in *Proceedings of the 2010 IEEE 30th International Conference on Distributed Computing Systems Workshops*, ser. ICDCSW '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 42–51. [Online]. Available: http://dx.doi.org/10.1109/ICDCSW.2010.35

[11] C. Basescu, A. Carpen-Amarie, C. Leordeanu, A. Costan, and G. Antoniu, "Managing data access on clouds: A generic framework for enforcing security policies," in *AINA*. IEEE Computer Society, 2011, pp. 459–466.

[12] M. K. Iskander, D. W. Wilkinson, A. J. Lee, and P. K. Chrysanthis, "Enforcing policy and data consistency of cloud transactions," in *Proceedings of the Second International Workshop on Security and Privacy in Cloud Computing*, ser. ICDCS-SPCC 2011. Washington, DC, USA: IEEE Computer Society, 2011.

[13] "CloudAudit: A6 - The Automated Audit, Assertion, Assessment, and Assurance API ," July 2011. [Online]. Available: http://cloudaudit.org

[14] "XEN networking blog," July 2011, http://wiki.xensource.com/xenwiki/XenNetworking.

[15] "VMware," July 2011. [Online]. Available: http://www.vmware.com

[16] N. M. K. Chowdhury and R. Boutaba, "A survey of network virtualization," *Computer Networks*, vol. 54, no. 5, pp. 862 – 876, 2010.

[17] "AMD Virtualization (AMD-V) Technology," July 2011. [Online]. Available: http://sites.amd.com/us/business/it-solutions/virtualization/Pages/amd-v.aspx

[18] "Intel virtualization," July 2011. [Online]. Available: http://www.intel.com/technology/virtualization/