Trust-based Resource Allocation for Cloud Datacenters: improving security and performance

Daniel Stefani Marcon, Miguel Cardoso Neves, Rodrigo Ruas Oliveira,
Luciana Salete Buriol, Luciano Paschoal Gaspary, Marinho Pilla Barcellos
Institute of Informatics – Federal University of Rio Grande do Sul
Email: {daniel.stefani, mcneves, ruas.oliveira, buriol, paschoal, marinho}@inf.ufrgs.br

***Abstract*—In this paper, we briefly revisit a resource allocation strategy we recently proposed that aims at mitigating the impact of selfish and malicious attacks in the intra-cloud network and at improving performance for applications. This is achieved by grouping applications from mutually trusting users into logically isolated domains (virtual infrastructures) with bandwidth guarantees. Evaluation results show the benefits of our strategy, which is able to offer better performance and network resource protection against attacks.**

I. Introduction

Cloud providers implement datacenters as highly multiplexed shared environments. However, they lack mechanisms to capture and control network requirements of the interactions among allocated virtual machines (VMs). For example, congestion control mechanisms used in intra-cloud networks do not ensure robust traffic isolation among different applications, especially with distinct bandwidth requirements [1]. This enables selfish and malicious use of the network, allowing tenants to launch performance interference (consumption of an unfair share of the network) and denial of service (DoS) attacks. These attacks result in unpredictable network performance for tenants. More specifically, when available bandwidth for an application goes below a certain threshold, the total application execution time is elongated monotonically [2]. Furthermore, the computation may have dependence on the data received from the network (if communication speed is reduced, the subsequent computation is delayed). Therefore, we improve network resource sharing among applications in the cloud by proposing a novel resource allocation scheme.

II. Resource Allocation Strategy

Our approach aims at improving both security and performance for tenant applications. The security is increased by mitigating the impact of selfish and malicious behavior in the intra-cloud network, while overall performance is improved by providing network performance isolation among applications. Unlike prior work [3], we investigate a strategy based on grouping of applications in logically isolated domains (virtual infrastructures - VIs) according to trust relationships between tenants, in order to maintain high resource utilization. Moreover, the proposed approach does not require any new hardware. In fact, it can be deployed either by configuring network devices or by modifying VM hypervisors. The reader may refer to [4] for details about the work.

III. Evaluation

To show the benefits of our approach in large-scale cloud platforms, we developed a simulator that models a multi-tenant shared datacenter (with 64,000 VMs) and implements our resource allocation strategy. We compare our strategy with a baseline scenario (current cloud allocation scheme), in which all tenants share the same network.

Increased security. Security is quantified by measuring the number of mutually untrusted tenants assigned to the same VI. It is desirable to have this value minimized, because it shows how exposed applications are to attacks. Figure 1 shows that the number of applications is not the main factor to increase security, but rather the number of VIs offered by the provider. We find that the number of mutually untrusted relationships decreases, and thereby security increases, logarithmically according to the number of VIs. Further, by reducing interference, applications achieve better performance, since interference in the network is one of the leading causes for poor application performance in the cloud [3].

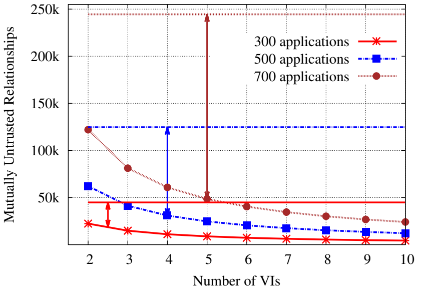

Fig. 1: Security when allocating applications in the cloud.

IV. Final Remarks

In summary, we have proposed a security- and network-performance-aware resource allocation strategy for Infrastructure as a Service (IaaS) cloud platforms, which can be applied on different datacenter network topologies, such as today's multi-rooted trees [3] and richer topologies (e.g., VL2).

Acknowledgements

This work has been supported by FP7/CNPq (Project SecFuNet, FP7-ICT-2011-EU-Brazil), RNP-CTIC (Project AltoStratus), as well as PRONEM/FAPERGS/CNPq (Project NPRV).

References

[1] A. Shieh, S. Kandula, A. Greenberg, C. Kim, and B. Saha, "Sharing the data center network," in *USENIX NSDI*, 2011.

[2] D. Xie, N. Ding, Y. C. Hu, and R. Kompella, "The only constant is change: Incorporating time-varying network reservations in data centers," in *ACM SIGCOMM*, 2012.

[3] H. Ballani, P. Costa, T. Karagiannis, and A. Rowstron, "Towards predictable datacenter networks," in *ACM SIGCOMM*, 2011.

[4] D. S. Marcon, R. R. Oliveira, M. C. Neves, L. S. Buriol, L. P. Gaspary, and M. P. Barcellos, "Trust-based grouping for cloud datacenters: improving security in shared infrastructures," in *IFIP TC6 Networking*, 2013.