

# AMNESiA: Affinity Measurement Platform for NFV-enabled Networks

Arthur Selle Jacobs, Ricardo Luis do Santos, Muriel Figueredo Franco, Eder John Scheid,  
Ricardo José Pfitscher, Lisandro Zambenedetti Granville  
Institute of Informatics — Federal University of Rio Grande do Sul  
{asjacobs, rlsantos, mffranco, ejscheid, rjpfitscher, granville}@inf.ufrgs.br

**Abstract**—AMNESiA is an affinity measurement platform for NFV-enabled networks, designed to consolidate and interpret existing monitoring data into an affinity metric, aiding operators to identify affinity and anti-affinity relations in the network. AMNESiA uses the latest snapshot of usage data, collected through a generic monitoring solution, from the database to measure affinity between VNFs.

## I. INTRODUCTION AND RELATED WORK

In Network Functions Virtualization (NFV)-enabled networks, inadequate placement of Virtualized Network Functions (VNFs) creates bottlenecks, impacting negatively on performance [1]. Network operators should establish affinity and anti-affinity rules between VNFs to avoid network and processing bottlenecks [2]. These rules can be based on several different aspects, such as VNFs minimum resource requirements, latency, or even network operators needs for the service [3]. Affinity has been addressed in the context of Cloud-based environments, in which solutions do not account for the functions being executed by each Virtual Machine (VM), focusing mainly on resource contention. In NFV-enabled networks, VNFs are chained in Forwarding Graphs (FGs) to provide services (*i.e.*, service chaining). Therefore, network operators must consider further than simply resource allocation when identifying affinity among VNFs. Most commonly, current solutions propose affinity relations based on CPU [4] and bandwidth [5]. In addition, most of these solutions do not propose any interaction or platform to aid operators on monitoring the state of the network.

Chen *et al.* [6] proposed a method to identify affinity relations based on resource demands and dependency among VMs. Franco *et al.* [7] presented VISION, a visualization platform with interactive and selective techniques for NFV-enabled networks. Network operators may use VISION to identify problems, determining affinity and anti-affinity relations between VNFs through visualizations. Yousaf and T. Taleb [8] proposed a fine-grained resource-aware VM management solution for NFV-enabled networks, based on a reference resource affinity score (RRAS). Their affinity calculation results on a vector quantity representing the impact of one reference Resource Unit (RU) (*e.g.*, memory) on other RUs (*e.g.*, processing, storage, and I/O module). Even though several authors have tackled issues regarding affinity, their solutions focus mainly on computational resources. Consequently, these solutions fall short for NFV-enabled networks, which require further

considerations when determining affinity. In addition to computational resources, requirements such as latency, FGs, and service performance, must be taken into account. Furthermore, most of current solutions lack in dynamicity, since they do not consider any interaction with operators, or simply do not provide a specific technique to identify affinity, relying solely on the expertise of the operator to do so.

In [9], we proposed an affinity measurement model for NFV-enabled networks based on an extensible set of criteria, coupled with input weights. The proposed affinity model may be used either by network operators to identify issues in the network, or by an orchestrator to resolve better placement for VNFs. Aiming to provide a Proof of Concept (PoC) for the former scenario, we developed AMNESiA, a platform that implements the proposed affinity measurement, aiding operators on monitoring the state of the network.

## II. AMNESiA

AMNESiA was built as a Web application, using Python, the Django Web Framework, and a PostgreSQL Database. AMNESiA is an affinity measurement platform, designed to consolidate and interpret existing monitoring data into an affinity metric, aiding operators to identify affinity and anti-affinity relations in the network. Therefore, we assume the data for VNFs, Physical Machines (PMs), and FGs are collected by any network monitoring solution, storing snapshots of monitored information into AMNESiA's database. AMNESiA then uses the latest snapshot of usage data to measure affinity between VNFs.

AMNESiA consists of three different features: (*i*) Affinity Measurement; (*ii*) Affinity Report, and (*iii*) Criteria Manager. Figure 1 depicts the AMNESiA architecture, as well as shows how AMNESiA integrates with a generic NFV environment. All three features developed in AMNESiA's Web-based interface use the same affinity measurement engine, which was implemented according to the affinity model introduced in [9]. Each AMNESiA feature is described in the next subsections.

### A. Affinity Measurement

The Affinity Measurement feature is a three-steps wizard to measure affinity between any two selected VNFs. On step one, the network operator has to select two VNFs from a list. We display on that list information regarding each VNF, such as

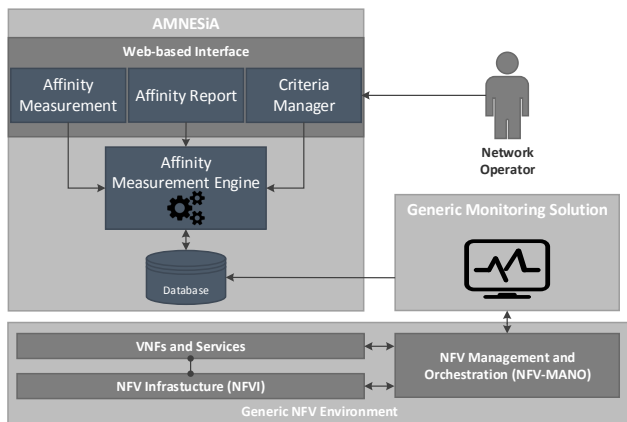


Figure 1: AMNESiA's Architecture.

name, PM on which the VNF is hosted on, usage percentages, and a list of FGs the VNF belong to.

Afterwards, the user has to provide weights to each affinity criterion. A network operator might give whatever weight he/she sees fit for each criterion. The weight of a criterion must be an integer larger or equal to zero. If an operator wants to disregard any criterion, he/she has to provide a weight value of zero for that criterion.

In the final step, the user is presented with the affinity values for the pair of selected VNFs, alongside the given criteria weights. As described on [9], if the selected VNFs are directly linked in more than one FG, the affinity measurement is calculated for each FG, since they may have different flows. If the VNFs are not chained in any FG, the affinity measurement is only calculated regarding PM criteria. Figure 2 presents the result of an affinity measurement between two VNFs that are chained in five FGs.

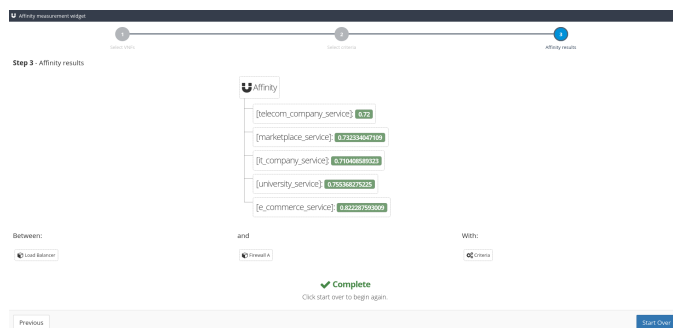


Figure 2: Affinity Measurement result screen.

### B. Affinity Report

The Affinity Report is also a three-step wizard. However, instead of selecting a pair of VNFs, the user selects an FG. The purpose is to measure the affinity of every two VNFs combination of an FG, providing an overview of the overall state of the service's affinity.

In this feature, the user must first select an FG from a list, in which we provide usage metrics regarding every flow of each FG, such as latency, bandwidth usage, and traffic in every link (either virtual or physical). Afterwards, the user must once again provide weights for each affinity criterion. After determining the weights of the criteria, the user is presented with the affinity results in the third step of the Affinity Report feature. As a result, we display a table containing all the two VNFs combination of the selected FG, alongside the affinity result for that pair. Since in this feature the user already provides an FG to measure affinity with, we only consider connections between VNFs regarding the selected FG.

### C. Criteria Manager

The Criteria Manager feature is designed to enable network operators to extend and customize the set of affinity criteria. We present network operators with a list of all the criteria in the application, which they might remove from the system according to their needs.

To add a new criterion, the user must fill a form with the new criterion's information, including criterion's name, type, and scope, the default weight, and the criterion's affinity equation. The criterion's affinity equation must be a Python function, receiving as input the objects of both VNFs being evaluated, the FG object to consider, and an object containing information from the Network Service Descriptor (NSD); and returning a numerical value between 0.001 and 1. After included, the new criterion can now be selected among existing criteria in other affinity measurement features.

### REFERENCES

- [1] Rashid Mijumbi, Joan Serrat, Juan-Luis Gorricho, Niels Bouten, Filip De Turck, and Raouf Boutaba. Network Function Virtualization: State-of-the-art and Research Challenges. *CoRR*, abs/1509.07675, 2015.
- [2] S. Oechsner and A. Ripke. Flexible support of VNF placement functions in OpenStack. In *Network Softwarization (NetSoft), 2015 1st IEEE Conference on*, pages 1–6, April 2015.
- [3] Alcatel-Lucent. Network Functions Virtualization - Challenges and Solutions. White paper, June 2013.
- [4] S. Sudevalayam and P. Kulkarni. Affinity-Aware Modeling of CPU Usage for Provisioning Virtualized Applications. In *Cloud Computing (CLOUD), 2011 IEEE International Conference on*, July 2011.
- [5] J. Sonnek, J. Greensky, R. Reutiman, and A. Chandra. Starling: Minimizing Communication Overhead in Virtualized Computing Platforms Using Decentralized Affinity-Aware Migration. In *2010 39th International Conference on Parallel Processing*, pages 228–237, Sept 2010.
- [6] J. Chen, K. Chiew, D. Ye, L. Zhu, and W. Chen. AAGA: Affinity-Aware Grouping for Allocation of Virtual Machines. In *Advanced Information Networking and Applications (AINA), 2013 IEEE 27th International Conference on*, pages 235–242, March 2013.
- [7] Muriel Franco, Ricardo Luis dos Santos, Alberto E. Schaeffer-Filho, and Lisandro Z Granville. VISION - Interactive and Selective Visualization for Management of NFV-Enabled Networks. In *The 30-th IEEE International Conference on Advanced Information Networking and Applications (AINA-2016)*, Centre de Congrès le Régent, Crans-Montana, Switzerland, March 2016.
- [8] F. Z. Yousaf and T. Taleb. Fine-grained resource-aware virtual network function management for 5G carrier cloud. *IEEE Network*, 30(2):110–115, March 2016.
- [9] Arthur Jacobs, Ricardo Luis dos Santos, Muriel Franco, Eder John Scheid, Ricardo José Pfitscher, and Lisandro Z Granville. Affinity Measurement for NFV-Enabled Networks: A Criteria-based Approach. In *15th IFIP/IEEE International Symposium on Integrated Network Management (IM 2017)*, May 2017. To appear.