NFV Anomaly Detection: Case-Study Through a Security Module

Lucas Bondan*, Tim Wauters†, Bruno Volckaert†, Filip De Turck†, Lisandro Zambenedetti Granville*
*UFRGS - INF, Brazil
†UGent - imec, Belgium

***Abstract*—Network Functions Virtualization (NFV) is a key networking concept whose benefits include scalability, flexibility, and cost-effective service provisioning. In NFV, Virtualized Network Functions (VNFs) are chained in Service Function Chains (SFCs) adaptable to customers' needs. VNFs and SFCs are sensitive elements that, if compromised, would affect network security. The detection of compromised VNFs and SFCs is imperative, and although anomaly detection can be used in such a context, there is a lack of research work on the use of anomaly detection in NFV. In this article, we exploit the use of anomaly detection mechanisms to identify suspicious VNFs and SFCs. We introduce, into a widely accepted NFV architecture, an NFV Security Module (NSM) that, by analyzing VNFs and SFCs' operations, detects anomalies possibly resulting from security attacks. To prove the concept, three mechanisms have been implemented and deployed in NSM to observe how anomaly detection performs, given quantitative and qualitative information. We found out that anomaly detection is effective for VNF and SFC security and, in the case of using entropy as anomaly detection technique, it presents accuracy of up to 98% without harming NFV environment operations.**

***Index Terms*—NFV, security, anomaly detection**

I. Introduction

Academia and industry have been exploiting Network Functions Virtualization (NFV), boosting innovation for network provisioning and management, and reducing Operational and Capital Expenditures (OPEX and CAPEX). As defined by the European Telecommunications Standards Institute (ETSI), NFV comprises the virtualization of functions originally performed by dedicated devices into software [1]. Such "softwarized" functions – called Virtualized Network Functions (VNFs) – are central to the NFV architecture. Virtualizing network functions through NFV brings flexibility to service delivery, given that customers' demands can be individually considered and dynamically adjusted through a chain of VNFs, composing Service Function Chains (SFC) (or VNF Forwarding Graphs - VNFFG, following ETSI nomenclature).

With the increasing deployment of NFV-enabled networks and NFV ecosystems consolidation, security-related issues started to gain attention [2]. Both virtualization and networking-related vulnerabilities are present in NFV environments, resulting in different types of threats (Figure 1). Also, undisclosed vulnerabilities (*i.e.,* zero-day threats) – constantly sought by security companies – enlarges the number of potential threats. Naturally, the consequences of an attacker exploiting NFV vulnerabilities can be devastating.

In NFV, observing anomalies considering both network and virtualized information helps identify threats and detect ongoing attacks. Therefore, compromised VNFs and SFCs need to be quickly detected, allowing network operators to apply suitable countermeasures, avoiding major harms to service delivery and customers' privacy. Considering the wide variety of NFV environments and the fact that anomaly detection mechanisms are appropriate tools to identify threats in different network contexts [3], the investigation of anomaly detection to increase NFV security would be expected. To the best of our knowledge, however, no other study has addressed the use of anomaly detection in VNFs and SFCs.

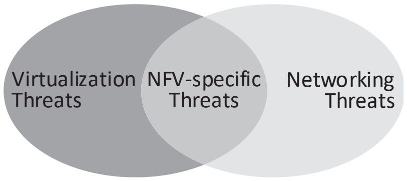

Figure 1. Threads affecting NFV environments [4].

In this article, we investigate the effectiveness of using anomaly detection mechanisms for NFV security, focusing on VNFs and SFCs. First, we provide an overeview of NFV adoption and evolution, highlighting aspects related to NFV security. We then revisit anomaly detection when employed in diverse network environments, which provides the foundation to use it in NFV. As a means to exploit anomaly detection in VNFs and SFCs, we then introduce an architectural framework called NFV Security Module (NSM). Then, taking into account realistic scenarios defined by ETSI, we present and discuss the characteristics of NFV threads, introduce the implementation of three anomaly detection mechanisms using NSM, and evaluate such mechanisms in the realistic scenarios. Finally, we close this article presenting conclusions and directions for future work.

II. Related Work

After the publication of the first NFV white paper, by ETSI, followed by its NFV architectural framework [5], different NFV initiatives from industry, academia, and standardization groups emerged. From industry, we highlight Telefónica OpenMANO, Cisco NFVI and AT&T ECOMP. At the time, companies were focused on realizing NFV as soon as possible, with security aspects addressed timidly.

Proposals like UNIFY, addressing service orchestration and automated service chaining; and T-NOVA, focusing on automated NFV Management and Orchestration (MANO), also emerged from academia. Later, FENDE was published [6], the first NFV marketplace and ecosystem with support to VNF distribution, execution, and SFC composition. Security only appeared in more recent works, such as the NFV security survey focused on 5G networks [7], defining a threat taxonomy specific for NFV-based 5G scenarios.

Open initiatives such as the Open Platform for NFV (OPNFV) and The Linux Foundation Open-O project started with the goal of developing open source NFV solutions. Later, ECOMP and Open-O merged to create ONAP, a platform for real-time VNF orchestration with a dedicated security coordination committee responsible for managing identified vulnerabilities and coordinating security-related activities. Despite the security concern, anomaly detection for NFV was not considered as a potential solution for NFV security.

From the standardization side, IETF established the SFC Working Group (SFCWG) and the NFV Research Group (NFVRG), both aiming at NFV-related challenges. ETSI announced NFV Security, identifying potential security vulnerabilities in NFV, making clear the importance of NFV security. Recently, ETSI released security enhancements for its NFV MANO architecture, considering communication-related security aspects. SFCWG released the latest version of its SFC protocol security draft, but once again without considering anomaly detection as an enabler for NFV security.

NFV solutions supported the development of security-related VNFs and SFCs, *i.e.,* NFV was a security enabler. However, no efforts towards securing NFV environments themselves – *i.e.,* NFV as the target of security – were observed. As such, despite the advantages of using anomaly detection for network security [8], anomaly detection was not considered for NFV environments security. We argue that anomaly detection is suitable for NFV environments security due to several reasons: (*i*) NFV offers a central control point of the network environment (*i.e.,* NFV orchestrator); (*ii*) NFV supports the easy collection of VNFs and SFC information; and (*iii*) NFV includes a dedicated MANO plane to enable automated actions. As such, the remainder of this article addresses this opportunity.

III. Anomaly Detection for Network Security

Anomalies are "*patterns in data that do not conform to expected behavior*" [3]. Such nonconforming patterns may result from problems in a system operation, *e.g.,* denial of service, and information leakage. The employment of anomaly detection in networking environments is based on the computation of a score that identifies the expected behavior of a monitored information or set of information. When such a score is not the one expected, it represents an anomaly. Anomalies detected in NFV environments can be results of events related to VNFs and SFCs, such as missing elements and misconfiguration. Such events can result from threats that, if exploited, may lead to services interruption and compromise the NFV environment. There exist four main anomaly detection technique groups [3]: (*i*) supervised training, (*ii*) statistical modeling, (*iii*) spectral theory, and (*iv*) information theory. Each group is more suitable for different information patterns and network environments.

Supervised training techniques require training datasets with regular behavior of the monitored system to enforce the proper training of the anomaly detection solution; they are often used to detect bulk anomalies in traffic flows. For example, reinforcement learning algorithms applied to identify malicious flows must know the regular behavior of the network to receive a reward or penalty after concluding an analysis, which will then be used to improve its knowledge about potential malicious activities [9]. However, given that VNFs and SFCs are deployed, migrated, and removed frequently, significant training datasets are unlikely to exist, thus preventing supervised training techniques as a viable option for NFV anomaly detection. Differently, statistical modeling techniques require precise characterization of both anomalous and regular behaviors. These techniques are often applied in intrusion detection systems where both regular and malicious behavior can be well characterized through mathematical models, which are hard to achieve in NFV [10]. Ultimately, both supervised training and statistical modeling miss the ability to detect potential unknown threats in NFV environments.

Despite their high accuracy, spectral theory techniques have high computational complexity and also need anomalous and regular instances to be separable in the lower dimensional embedding of the data, *i.e.,* the variation between normal and anomalous data must be high enough to separate them. Spectral theory techniques have been applied to network intrusion detection systems with high computational capacity, such as cloud computing systems where dedicated servers can be employed to execute the spectral theory-based algorithms. Information theory techniques, however, do not require training datasets or statistical models and present less complexity than spectral theory techniques, demanding fewer resources to run in acceptable time. For instance, entropy-based techniques can be employed in different environments, requiring only the set of monitored information to characterize the network environment, using it as baseline for further anomaly detection analyses [8]. Such characteristics turns information theory a strong candidate to be employed in NFV environments, with two main requirements: (*i*) **wide view of the NFV environment**, since anomaly detection mechanisms require information regarding all NFV elements being monitored; and (*ii*) **non-blocking information access**, since anomaly detection runs in parallel with the NFV environment operation.

Motivated by (*i*) the lack of solutions to cover the security attributes presented, (*ii*) NFV environments potential vulnerabilities, and (*iii*) the valuable results of anomaly detection mechanisms, we first introduced an architectural framework that allows designing and implementing anomaly detection mechanisms for NFV environments [11]. Now, we advance our previous investigation by (*i*) improving the proposed NFV Security Module (NSM) architecture to better fit in different scenarios, (*ii*) using realistic evaluation scenarios based on ETSI definitions, (*iii*) adding new anomaly detection mechanisms to improve detection accuracy considering different types of data, and (*iv*) extending the evaluation performed.

IV. Providing Anomaly Detection Capabilities to NFV Environments

NFV Security Module (NSM) extends ETSI's NFV architecture to support anomaly detection. The lower portion of Figure 2 depicts ETSI's NFV original elements, with an NFV Infrastructure (NFVI) composed of physical and virtual resources (*e.g.,* memory, CPU, network). Such resources are consumed by VNFs and each running VNF is managed by an Element Management System (EMS).

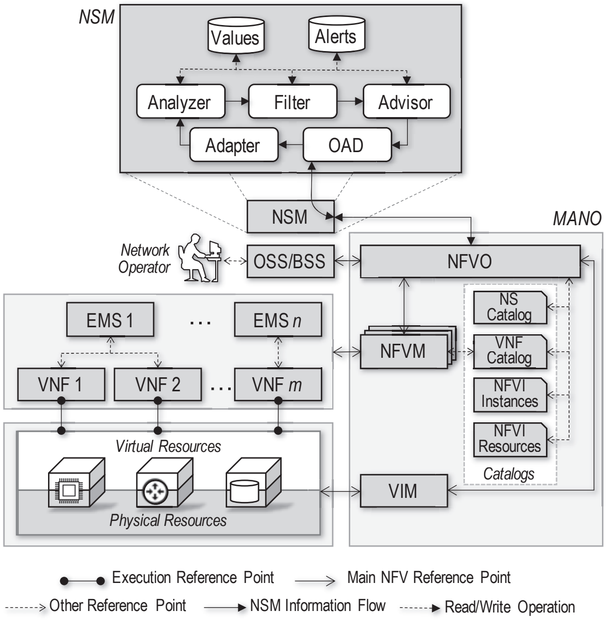

Figure 2. Detailed NSM architectural framework integrated with ETSI NFV.

MANO (Figure 2, bottom-right) interacts with physical and virtual resources via Virtual Infrastructure Manager (VIM), and with EMSes via VNF Managers (VNFMs). While VIM is responsible for resources management, VNFMs are responsible for VNF and SFC life-cycle management. NFV Orchestrator (NFVO) orchestrates the NFV environment interacting with VNFMs and VIM. Catalogs stores important information about Network Services, VNF images, NVFI, and NFV instances. The network operator manages VNFs and SFCs through Operations and Business Support Systems (OSS/BSS) that interact with NFVO.

NSM (Figure 2, top) communicates with NFVO to: (*i*) retrieve VNFs and SFCs information needed in anomaly detection analysis, and (*ii*) notify NFVO once anomalies are detected. An anomaly detection analysis is triggered either when an NSM internal interval expires or whenever NFVO requests an analysis, *e.g.,* when NFVO touches VNFs or SFCs. NSM was designed considering: (*i*) running VNFs and SFCs information acquired by the NFVO (*monitored information*), and (*ii*) information stored in the catalogs defined by ETSI and managed by the NFVO (*cataloged information*).

Orchestrator Abstraction Driver (OAD) creates an abstraction layer hiding from other NSM components the specificities of different NFVOs. OAD retrieves VNF and SFC information and forwards it to the **Adapter**. The Adapter then converts it into a format suitable to the anomaly detection mechanism in the Analyzer considering its implementation.

Anomaly detection is performed in the **Analyzer** using information received from the Adapter and information available in the Values and Alerts databases. If no anomaly is detected, the Analyzer ends the analysis and updates the Values database. However, if an anomaly is detected, the Analyzer forwards the detection information to the Filter.

The **Filter** identifies whether anomalies reported by the Analyzer are threats. If an anomaly results from a legitimate behavior, the Values database is updated and the analysis ends. Otherwise, if a threat is identified, Filter forwards the associated information to the **Advisor**. If Filter is unable to determine whether an anomaly is a threat, then the anomaly is classified as resulting from a *potential threat* and that information is forwarded to Advisor.

For each detected (potential) threat, the Advisor, by using recommendation algorithms, computes mitigation actions to be suggested to NFVO. The suggested actions seek to mitigate potential threats or known attacks that may be related to the threat occurrence. Then, the Advisor issues alerts composed of: (*i*) identified (potential) threats, (*ii*) affected VNFs and SFCs, and (*iii*) suggested actions. Such alerts are recorded at the Alerts database and forwarded to NFVO, who will decide upon executing or not the suggested actions.

Values and Alerts offers stored information back to the other components to further improve anomaly and threat detection. For example, the Analyzer can use the informa-tion stored in Values to learn and enhance future analysis. These communications close the interactions between the NSM internal components. Still, NSM and NFVOs oper-ate together to seek VNFs integrity, availability, and confi-dentiality. For further details regarding NSM, please access *lume.ufrgs.br/bitstream/handle/10183/197460/001097713.pdf* .

V. Case Study

Our case study is based on an environment where an NFVI provider manages sets of VNFs and SFCs owned both by that provider itself as well as by 3rd party virtual network service providers. These 3rd party providers rent the NFVI provider's infrastructure to host some (or all) of their VNFs and SFCs. This scenario is referred to as "hosted virtual network operators" by ETSI NFV-SEC [4]. In this scenario, VNFs can belong to both the NFVI operator or customers, but regardless of the VNF owner, all SFC are handled by the NFVI operator through NFV MANO, as depicted in Figure 3.

Figure 3 depicts three different SFCs, along with their paths through the servers of the operator's NFVI. VNFs can belong to both the NFVI operator or customers, but regardless of the VNF owner, all SFC are handled by the NFVI operator through the NFV MANO plane. SFC 1 has its endpoint inside the NFVI, which indicated the service is consumed by the NFVI provider itself, *e.g.,* performing predictive caching for content delivery networks. In this case, another SFC can be instantiated to deliver the service to a given customer when requested. To carry out our evaluation, we also assume that:

- VNFs and SFCs are free of bugs;

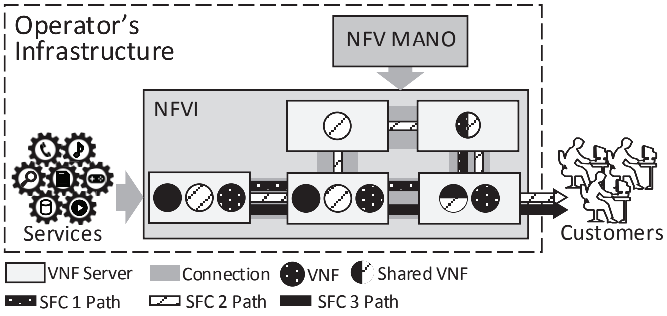

Figure 3. Example of the hosted virtual network operators scenario.

- Network connections are stable;
- No human intervention during the anomaly detection process; and
- NVFI is able to accommodate all VNF and SFC demands.

The following aspects are outside the scope of this study:

- Threats related to human error in network operation;
- VNFs and SFCs verification and validation;
- Infrastructure errors (network and resources);
- Dependability attributes: reliability and maintainability.

A. Anomaly Detection Mechanisms

The anomaly detection mechanisms designed and implemented use both monitored and cataloged information during the analysis, calculating both monitored information entropy and cataloged information entropy. Different Shannon's Information Entropy-based anomaly detection mechanisms were designed and implemented, considering two types of data: (*i*) *qualitative*, values interpreted as properties and attributes (*i.e.,* qualities), such as VNF identifiers; and (*ii*) *quantitative*, values numerically analyzed and processed, such as the customers' bandwidth. Both information types are analyzed separately. While small variations in quantitative information may occur and not indicate an anomaly, tiny variations in qualitative information may indicate an inconsistency and a potential threat. This way, one quantitative detector (*Numerical Entropy-based Detector* - NED) and two qualitative detectors (*Single Entropy-based Detector* - SED and *Merged Entropy-based Detector* - MED) were implemented. The main reasons to use Shannon's Information Entropy are its low complexity ($O(n)$), resulting in low impact in the NFV environment; and its wide dissemination and adoption across different research areas [3].

A previous investigation showed that SED has fast execution time and proven effectiveness [11] . However, SED may present false-negatives when, *e.g.,* the amount of missing elements matches the amount of unregistered elements in the monitored information. Such a situation may cause SED's monitored information entropy to remain unchanged in comparison with SED's cataloged information entropy, even when anomalies are occurring. MED was designed to avoid such false-negative detection, refining the qualitative information entropy calculation by merging cataloged and monitored information into a merged list. As such, if an unregistered or missing element occurs in the monitored information, merged information entropy will differ from cataloged information entropy, indicating an anomaly. An example of a missing element is a VNF identifier not present in the monitored information of an SFC, but present in the cataloged information. Similarly, an unregistered element could be an additional port in the VNF monitored information that is not present in the cataloged information of that VNF.

For quantitative information, merging monitored and cataloged information is not mandatory, due to their discrete nature which makes it virtually impossible same entropy variations to appear in two different analyses. However, distinct entropy results may appear from every evaluation, turning entropy-to-entropy comparison ineffective. NED was designed to surpass this issue, analyzing the monitored information entropy considering historical cataloged information entropies. Anomalies for quantitative information are detected by analyzing whether the monitored information entropy fits into the interval composed of the mean of historical entropy values plus/minus its standard deviation. In addition, a parameter (β) is defined to adjust the interval size. The higher the value of β, the more permissive the detector is, with $\beta = 1$ representing no changes. The mechanisms receive as parameters the input data resulting from an adapted algorithm (described in detail in Subsection V-C) and, in the case of NED, a β value. As such, it is imperative for datasets to be consistent, since the accuracy of the anomaly detection relies on the integrity of the data used as input. The detectors' Python implementation is available at

Table I
ANOMALIES AND THREAT CHARACTERISTICS

Threat	Potential Attack	Risk	Security Attribute
Missing SFC element	DoS	Service stops working or not working properly	Availability
Unauthorized bandwidth allocation	Privilege escalation	Users receive privileges above stipulated, network traffic congestion	Integrity
Uncataloged/modified VNF	Man-in-the-middle	Unauthorized users access VNFs and SFCs, information leakage	Confidentiality
	Privilege escalation	Users receive privileges above stipulated, network congestion	Integrity
Uncataloged/modified connection point and virtual link	Man-in-the-middle	Information leakage to unauthorized users or attackers	Confidentiality
	Privilege escalation	Users receive privileges above stipulated	Integrity

github.com/ComputerNetworks-UFRGS/nsm.

B. Threats, Potential Attacks, and Risks

An anomaly results from the occurrence or change of a particular set of circumstances, *i.e.,* from an *event*. When a new or newly discovered event has the potential to harm a system, it represents a *threat* [12]. A threat may be the result of different *attacks*, representing a *risk* to the NFV environment. Table I presents the threats, potential attacks, risks, and related security attributes considered in the case study. We have selected those threats because they can affect VNFs and SFCs in different scenarios. Threats are detected at both VNF and SFC levels. Depending on the type of threat detected, it is possible to generate a cascade effect on other VNFs part of the same SFC. This cascade effect can also be detected since it will generate even greater anomalies.

The threat "Missing SFC element" may indicate a DoS attack, compromising the availability of NFV service provisioning. "Unauthorized bandwidth allocation" may indicate a potential attack of privilege escalation, associated with the risk of users receiving privileges above stipulated and compromising the integrity security attribute. "Unauthorized/modified VNF" may indicate both Man-in-the-middle and privilege escalation attacks. While Man-in-the-middle represents risks of unauthorized users accessing VNFs and SFCs and information leakage – compromising the confidentiality security attribute –, the privilege escalation attack may indicate risks of users receiving privileges above stipulated and network congestion – an integrity break. "Uncatalogued/modified connection point and virtual link" may be related to multiple attacks too, *i.e.,* Man-in-the-middle and privilege escalation. While the first attack may represent a risk of information leakage to unauthorized users or attackers – compromising confidentiality –, the second has a potential risk of users receiving privileges above stipulated, compromising the integrity security attribute.

Considering that every non-conforming pattern detected on VNF and SFC cataloged/monitored information results in an anomaly that may be related to a threat, every operation related to VNF and SFC is considered by the proposed solutions. In summary, our solutions are agnostic regarding the operation that generates the anomaly.

C. Evaluation

An algorithm was designed to create the datasets that were analyzed by NSM in our experiments, adapting the algorithm of Rankothge *et al.* [13] to operate as follows. The algorithm receives as input: (*i*) the average number of SFCs (defined as 100, reflecting large scale enterprise networks, where each SFC is composed of 2 to 7 VNFs [14]); (*ii*) the average number of VNFs (the number of VNFs for a given customer considered follows a truncated power-low distribution with exponent 2, minimum 2, and maximum 7 [13]); (*iii*) the threat event likelihood (defined as 60% based on enterprise reports [15]); and (*iv*) the legitimate event likelihood (also 60%). Both events likelihood follow a normal distribution.

The algorithm generates snapshots composed of two datasets: (*i*) monitored data and (*ii*) cataloged data. Each snapshot contains a timestamp, and each line of the datasets contains: SFC identifier, V NFs c omposing t he S FC, connection points, virtual links, and user's allocated bandwidth. Whenever a new event (legitimate or threat) occurs, a new snapshot is generated. Legitimate events can be the registration of a new SFC or VNF, reallocation of users' bandwidth, VNFs re-routing within an SFC (*i.e.,* changes in connection points of virtual links), among others. In turn, threats are represented by events related to Table I. When NSM operation starts (triggered by an event or a time interval), it considers the most recent datasets.

Our analysis compares NSM detection results with generated datasets, observing the accuracy of the detection mechanisms and the detection time of each trigger. We argue that accuracy and detection time are the most important outcomes when it comes to anomaly detection. Other parameters such as execution time and resource consumption could also be analyzed, but since these parameters may be affected by the network and hardware employed, accuracy was chosen to prove the effectiveness of the proposed mechanisms. For the analysis of the detection time of each trigger, a single detection mechanism is used (MED), since the execution time of the different mechanisms implemented is negligible when compared to the detection time of the triggers.

1) Detection Accuracy: False Positive Rate (FPR) and True Positive Rate (TPR) are considered for this evaluation. NED starts its analysis with 100 entropy values available at the Values database, as mentioned in Subsection V-A, and three values of β are considered: 1 (no change), 0.5 (half-size), and 2 (double-size). Results are presented in Figure 4 using a Receiver Operating Characteristic curve (ROC curve), which shows TPR on the y-axis and FPR on the x-axis. The closer to the top-left corner the higher the detector accuracy. As a baseline for comparison, a random detection line is presented together with detectors' results.

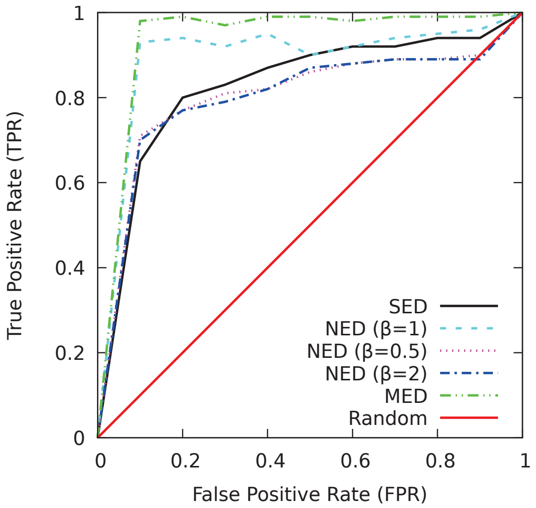

Figure 4. Accuracy of SED, MED, and NED with three different β values.

MED presents higher accuracy in all cases (around 98%), followed by NED with $\beta = 1$, SED, and NED with $\beta = 2$ and $\beta = 0.5$. MED great results can be ascribed to the merged list composition, which minimizes false-negative occurrences by calculating the merged information entropy.

With 95% on average, NED accuracy is slightly smaller than MED, which can be assigned to scenarios where small variation in the monitored allocated bandwidth may not be detected, especially at the start of NED execution, when few monitored values are available to compute the mean and standard deviation to characterize the monitored elements entropy. Thus, NED can present a high number of false positives until it has a certain number of samples. Afterwards, the tendency is for the number of false positives to fall.

NED accuracy decreased using a bigger β (2), because of the higher tolerance when using a higher β. It means that NED might consider greater changes in the entropy as normal when they might be anomalies. Still, using a smaller β may restrict too much the analyzed samples, and regular information might be considered anomalies too.

2) Detection Time: considers two NSM detection triggers: NSM internal interval and NFVO analysis requests. The first trigger is referred to as *Interval*, while the second one as *Request*. A period of 60 minutes was considered, using MED because of its higher accuracy for qualitative information.

In the *Request* analysis, NSM is configured to execute whenever a new legitimate event occurs, *e.g.,* new VNFs registration or configuration changes in existing VNFs. Such events occurrences (regular or anomalous) varies from 1 to 10 per hour, following the 60% likelihood defined. In the *Interval* analysis, intervals to analyze the monitored information are configured in NSM from 1 to 60 minutes. Figure 5 depicts the results obtained.

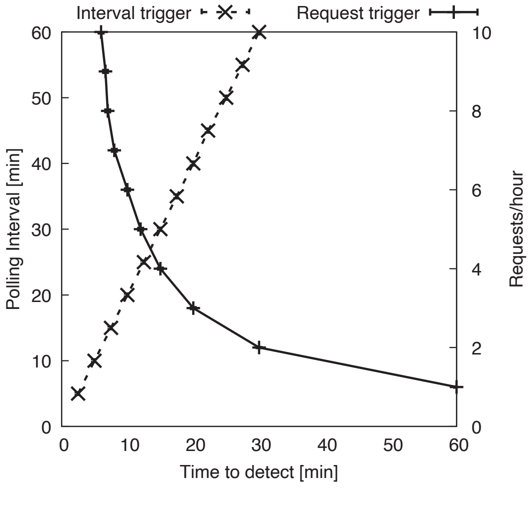

Figure 5. Detection time comparison of both triggers.

As the trigger interval increases, the anomaly detection delay increases linearly considering the *Interval* trigger. With the *Interval* trigger, the detection time takes half of the polling interval to detect anomalies, on average. In turn, the *Request* trigger presents decreasing logarithmic time to detect anomalies as the number of notifications per hour increases linearly. There is an intersection point where both *Request* and *Interval* triggers present the same detection time, (12 min), when a polling interval of 24 min and 5 notifications per hour are used. In highly dynamic scenarios where VNFs and SFCs information changes often, using the *Request* trigger activates NSM more often, so anomalies might be detected faster without performing unnecessary analysis, which may occur when using short intervals for the *Interval* trigger.

The effectiveness of the *Interval* trigger is directly related to the interval configured. Short intervals may detect anomalies faster, but imply in more NSM executions. In turn, higher intervals imply fewer NSM executions, but anomalies will take longer to be detected. The *Request* trigger can be configured with different strategies, such as executing the anomaly detection whenever the NFVO acquires information from VNFs and SFCs (monitoring events). However, processing time of both NSM and NFVO may increase with such a strategy, and NFVOs should support operations parallelism over monitored information, *i.e.,* monitored information acquisition, forward it to NSM, receive back the anomaly detection results, and evaluate the application or not of the suggested actions.

VI. Conclusion and Future Research

This article discussed the advancements related to NFV environments' security, from its definition to recent proposals regarding NFV in emerging network environments. An NFV Security Module (NSM) was presented to investigate anomaly detection effectiveness for NFV security. We analyzed if threats related to security attributes could be properly detected using anomaly detection, which led to the design, implementation, and evaluation of three different entropy-based anomaly detection mechanisms. A case study with a realistic NFV scenario was considered for our experiments, allowing us to conclude that anomaly detection effectively identifies potential threats in NFV environments, presenting accuracy of up to 98% among the entropy-based mechanisms designed. Also, two detection triggers were analyzed (*Request* and *Interval*), presenting both linear and logarithmic detection times depending on the trigger and configuration used. As future research, new mechanisms can be designed and evaluated using NSM, considering real-time resource consumption by container engines and virtual machines.

References

[1] R. Mijumbi, J. Serrat, J.-L. Gorricho, N. Bouten, F. De Turck, and R. Boutaba, "Network Function Virtualization: State-of-the-Art and Research Challenges," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 236–262, Mar. 2016.

[2] W. Yang and C. Fung, "A Survey on Security in Network Functions Virtualization," in *PIEEE NetSoft Conference and Workshops (NetSoft)*. Seoul, South Korea: IEEE, Jun. 2016, pp. 15–19.

[3] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly Detection: A Survey," *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, Jul. 2009.

[4] B. Briscoe *et al.*, "Network Functions Virtualisation (NFV) - NFV Security: Problem Statement," ETSI NFV ISG, Online, Tech. Rep., 2014, available at: https://www.etsi.org/deliver/etsi_gs/NFV-SEC/001_099/001/01.01.01_60/gs_NFV-SEC001v010101p.pdf. Accessed on Aug. 2021.

[5] M. Chiosi *et al.*, "Network Functions Virtualisation (NFV) - Use Cases," ETSI NFV ISG, Online, Tech. Rep., 2013, available at: https://www.etsi.org/deliver/etsi_gs/nfv/001_099/001/01.01.01_60/gs_nfv001v010101p.pdf. Accessed on Apr., 2021.

[6] L. Bondan, M. F. Franco, L. Marcuzzo, G. Venancio, R. L. Santos, R. J. Pfitscher, E. J. Scheid, B. Stiller, F. De Turck, E. P. Duarte, A. E. Schaeffer-Filho, C. R. P. d. Santos, and L. Z. Granville, "FENDE: Marketplace-Based Distribution, Execution, and Life Cycle Management of VNFs," *IEEE Communications Magazine*, vol. 57, no. 1, pp. 13–19, Jan. 2019.

[7] T. Madi, H. A. Alameddine, M. Pourzandi, and A. Boukhtouta, "NFV security survey in 5G networks: A three-dimensional threat taxonomy," *Computer Networks*, vol. 197, 2021.

[8] A. Santos da Silva, J. A. Wickboldt, L. Z. Granville, and A. Schaeffer-Filho, "ATLANTIC: A framework for anomaly traffic detection, classification, and mitigation in SDN," in *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Apr. 2016, pp. 27–35.

[9] M. Yousefi, N. Mtetwa, Y. Zhang, and H. Tianfield, "A Reinforcement Learning Approach for Attack Graph Analysis," in *IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*, Aug. 2018, pp. 212–217.

[10] M. A. Ahmed and Y. A. Mohamed, "Enhancing Intrusion Detection Using Statistical Functions," in *2018 International Conference on Computer, Control, Electrical, and Electronics Engineering (ICCCEEE)*, Aug. 2018, pp. 1–6.

[11] L. Bondan, T. Wauters, B. Volckaert, F. D. Turck, and L. Z. Granville, "Anomaly Detection Framework for SFC Integrity in NFV Environments," in *IEEE Conference on Network Softwarization (NetSoft)*. Bologna, Italy: IEEE, Jul. 2017, pp. 1–5.

[12] "ISO/IEC 27001: Information technology - Security Techniques - Information Security Management Systems," International Organization for Standardization, Geneva, CH, Standard, 2013.

[13] W. Rankothge, F. Le, A. Russo, and J. Lobo, "Data Modelling for the Evaluation of Virtualized Network Functions Resource Allocation Algorithms," *Computing Research Repository (CoRR)*, 2017, available at: http://arxiv.org/abs/1702.00369. Accessed on Aug. 2021.

[14] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar, "Making Middleboxes Someone else's Problem: Network Processing As a Cloud Service," in *ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*. Helsinki, Finland: ACM, 2012, pp. 13–24.

[15] D. Anstee, P. Bowen, C. Chui, and G. Sockrider, "Worldwide infrastructure security report," Arbor Networks, White Paper, Tech. Rep., 2017, available at: https://www.arbornetworks.com/insight-into-the-global-threat-landscape. Accessed on Jan. 2021.

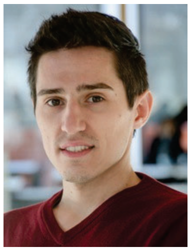

Lucas Bondan is an R&D Coordinator at the Brazilian National Research and Educational Network (RNP) and Supervisor Professor at University of Brasília. His research interests include NFV security, management & orchestration, and SFC.

Tim Wauters is a post-doctoral fellow at Ghent University. His research interests include network and service architectures and management solutions for scalable multimedia delivery services.

Bruno Volckaert is a professor at Ghent University in Belgium. His research interests include Cloud computing advances, NFV/SFC (a.o. IETF standardisation), application level distributed system design and architecture, and distributed (data) management systems for Smart City/Smart Transportation.

Filip De Turck is a Full Professor at Ghent University in Belgium, where he leads the network and service management research group. His research interests include scalable software architectures for network and service management, design and performance evaluation of novel QoE-aware multimedia delivery systems.

Lisandro Zambenetetti Granville is a Full Professor at the Federal University of Rio Grande do Sul (UFRGS) in Brazil. His research interests include management of network virtualization, intent-based networking, SDN, NFV, and network programmability.