NFV Anomaly Detection: Case Study through a Security Module

Lucas Bondan, Tim Wauter, Bruno Volckaert, Filip De Turck, and Lisandro Zambenedetti Granville

The authors exploit the use of anomaly detection mechanisms to identify suspicious VNFs and SFCs. They 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.

Abstract

Network function virtualization (NFV) is a key networking concept whose benefits include scalability, flexibility, and cost-effective service provisioning. In NFV, service function chains (SFCs) adaptable to customers' needs are created by chaining virtualized network functions (VNFs). 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 percent without harming NFV environment operations.

Introduction

Network function virtualization (NFV) has been explored by both industry and academia, boosting innovation for network provisioning and management, and reducing operational and capital expenditures (OPEX and CAPEX). According to the European Telecommunications Standards Institute (ETSI) definition, 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, called service function chain (SFC) (or VNF forwarding graph, following ETSI nomenclature).

With the increasing deployment of NFV-enabled networks and NFV ecosystems consolidation, NFV security started to be explored [2]. Both virtualization and networking-related vulnerabilities are present in NFV environments, resulting in different types of threats (Fig. 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. Taking in to account 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.

In this article, we investigate the effectiveness of using anomaly detection mechanisms for NFV security, focusing on VNFs and SFCs. First, we provide an overview 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 threats, 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.

Related Work

After ETSI released the first NFV white paper, followed by its NFV architectural framework [4], 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

Digital Object Identifier:
10.1109/MCOM.001.2100408

Lucas Bondan and Lisandro Zambenedetti Granville are with UFRGS — INF, Brazil;
Tim Wauters, Bruno Volckaert, and Filip De Turck are with UGent — IMEC, Belgium.

0163-6804/22/$25.00 © 2022 IEEE

T-NOVA, focused on automated NFV management and orchestration (MANO), also emerged from academia. Later, FENDE was published [5], the first NFV marketplace and ecosystem with support for VNF distribution, execution, and SFC composition. Security only appeared in more recent works, such as the NFV security survey focused on 5G networks [6], defining a threat taxonomy specifically 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 the real-time VNF orchestration platform ONAP, with a dedicated security coordination committee responsible for coordinating security-related activities such as managing identified vulnerabilities. Despite the security concern, anomaly detection for NFV was not considered as a potential solution for NFV security.

From the standardization side, the Internet Engineering Task Force (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 again not 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 toward 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 [3], anomaly detection was not considered for NFV environments' security. We argue that anomaly detection is suitable for NFV environments' security due to several reasons, including:

1. NFV offers a central control point of the network environment (i.e., NFV orchestrator).
2. NFV supports the easy collection of VNFs and SFC information.
3. NFV includes a dedicated MANO plane to enable automated actions.

Thus, the remainder of this article addresses this opportunity.

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 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 monitored information or a set of information. When such a score is not the one expected, it represents an anomaly. Anomalies detected in NFV environments can be as a result 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 service interruption and compromise the NFV environment. There are four main anomaly detection technique groups [3]:

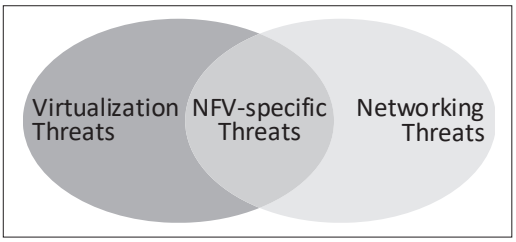

FIGURE 1. Threats affecting NFV environments [11].

1. Supervised training
2. Statistical modeling
3. Spectral theory
4. 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 [7]. 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. In comparison, 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 [8]. 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; that is, normal and anomalous data must present a variation high enough to be separated. 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 spectral-theory-based algorithms. In comparison, neither training datasets nor statistical models are required by information theory techniques, which also 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 a baseline for further anomaly detection analyses [3]. Such characteristics make information theory a strong candidate to be employed in NFV environments, with two main requirements:

1. **Wide view of the NFV environment**, since anomaly detection mechanisms require information regarding all NFV elements being monitored

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.

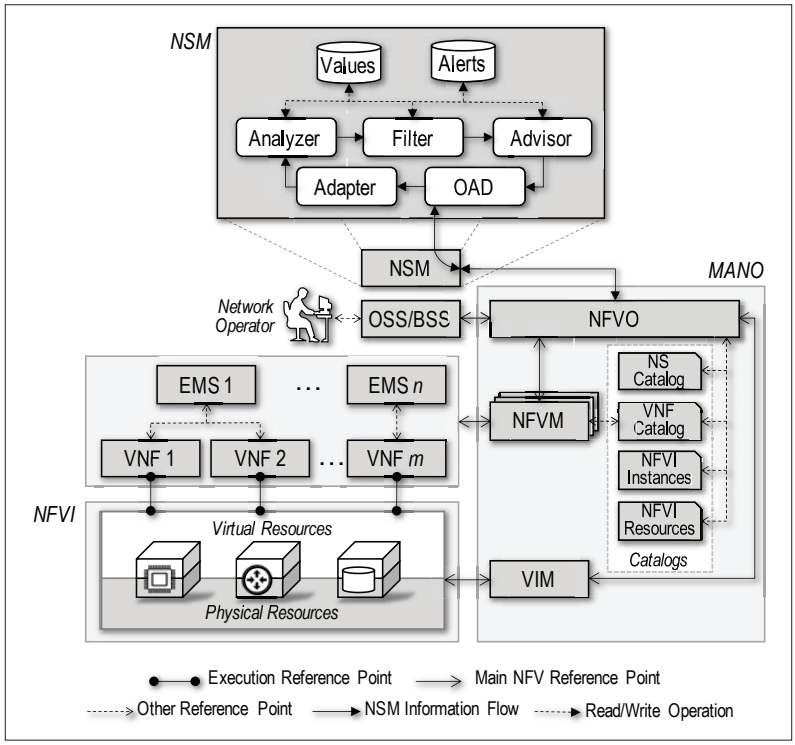

FIGURE 2. Detailed NSM architectural framework integrated with ETSI NFV.

2. **Non-blocking information access**, since anomaly detection runs in parallel with the NFV environment operation.

Motivated by the lack of solutions to cover the security attributes presented, NFV environments' potential vulnerabilities, and anomaly detection valuable results, we first introduced an architectural framework that allows designing and implementing anomaly detection mechanisms for NFV environments [9]. Now, we advance our previous investigation by:

1. Improving the proposed NFV Security Module (NSM) architecture to better fit in different scenarios
2. Using realistic evaluation scenarios based on ETSI definitions
3. Adding new anomaly detection mechanisms to improve detection accuracy considering different types of data
4. Extending the evaluation performed

Providing Anomaly Detection Capabilities to NFV Environments

The NFV security module (NSM) extends ETSI's NFV architecture to support anomaly detection. The lower portion of Fig. 2 depicts ETSI's NFV original elements, with an NFV infrastructure (NFVI) composed of physical and virtual resources (e.g., memory, CPU, network). These resources are consumed by VNFs, and each running VNF is managed by an element management system (EMS).

MANO (Fig. 2, bottom right) interacts with physical and virtual resources via the virtual infrastructure manager (VIM), and with EMSs via VNF managers (VNFMs). While the VIM is responsible for resource management, VNFMs are responsible for VNF and SFC life cycle management. The NFV orchestrator (NFVO) orchestrates the NFV environment interacting with VNFMs and VIM. Catalogues store important information about network services, VNF images, NVFI, and NFV instances. The network operator manages VNFs and SFCs using operations and business support systems (OSS/BSS) that interact with the NFVO.

NSM (Fig. 2, top) communicates with the NFVO to:

1. Retrieve VNFs and SFCs information needed in anomaly detection analysis
2. Notify the NFVO when anomalies are detected

An anomaly detection analysis is triggered either when an NSM internal interval expires or whenever the NFVO requests an analysis (e.g., when the NFVO touches VNFs or SFCs). NSM was designed considering:

1. Running VNFs and SFCs information acquired by the NFVO (*monitored information*)
2. Information stored in the catalogs defined by ETSI and managed by the NFVO (*catalogued information*)

The **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 for 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 legitimate behavior, the Values database is updated and the analysis ends. Otherwise, if a threat is identified, the Filter forwards the associated information to the **Advisor**. If the Filter is unable to determine whether an anomaly is a threat, the anomaly is classified as resulting from a *potential threat*, and that information is forwarded to the Advisor.

For each detected (potential) threat, the Advisor, by using recommendation algorithms, computes mitigation actions to suggest 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:

1. Identified (potential) threats
2 Affected VNFs and SFCs
3. Suggested actions

Such alerts are recorded at the Alerts database and forwarded to NFVO, who will decide whether or not to execute the suggested actions.

Values and Alerts offer stored information back to the other components to further improve anomaly and threat detection. For example, the Analyzer can use the information stored in Values to learn and enhance future analysis. These communications close the interactions between the NSM internal components. Still, NSM and NFVOs operate together to seek VNFs' integrity, availabil-

ity, and confidentiality. Further details regarding NSM can be found in [10].

Case Study

Our case study is based on an environment where an NFVI provider manages sets of VNFs and SFCs owned by both that provider itself as well as third-party virtual network service providers. These third-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 [11]. In this scenario, both the NFVI operator and customers own VNFs, although all SFCs are controlled by the NFVI operator using NFV MANO regardless of the VNFs' owner, as depicted in Fig. 3.

In Fig. 3, three SFCs and their specific paths through the operator's NFVI servers are depicted. For instance, SFC 1's endpoint is inside the NFVI, indicating a service consumed by the NFVI provider itself (e.g., predictive caching for content delivery networks), so an additional SFC can be created 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.
- Network connections are stable.
- There is no human intervention during the anomaly detection process.
- NVFI is able to accommodate all VNF and SFC demands.

The following aspects are outside the scope of this study:

- VNF and SFC validation and verification
- Threats associated with human error in network operation
- Infrastructure errors (resources and network)
- Dependability attributes (i.e., maintainability and reliability)

Anomaly Detection Mechanisms

The anomaly detection mechanisms designed and implemented use both monitored and catalogued information during the analysis, calculating both monitored information entropy and catalogued information entropy. Different Shannon's information-entropy-based anomaly detection mechanisms were designed and implemented, considering two types of data:

1 *Qualitative*, values interpreted as properties and attributes (i.e., qualities), such as VNF identifiers

2 *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 could indicate an inconsistency related to 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

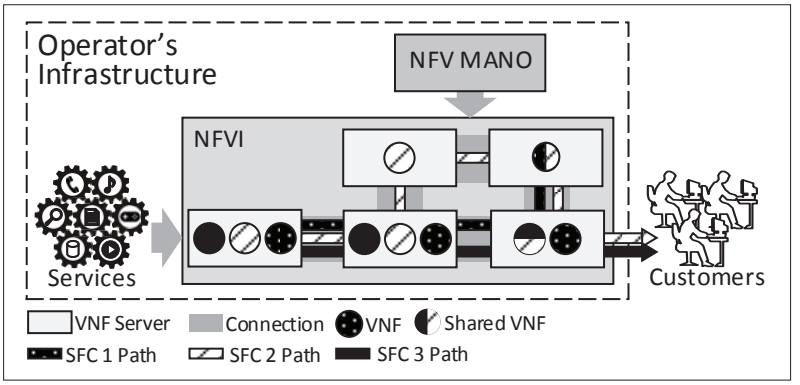

FIGURE 3. Hosted virtual network operators scenario example with three SFCs.

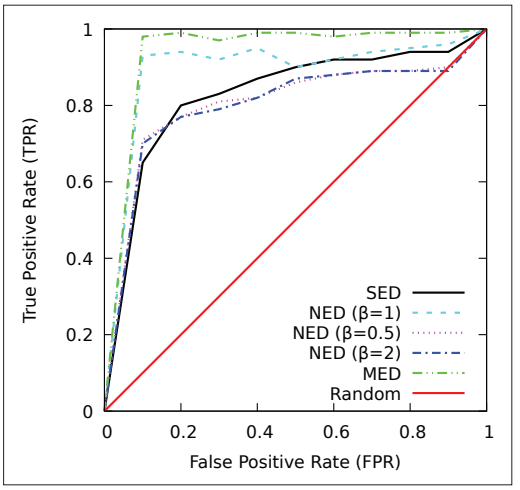

FIGURE 4. Accuracy of SED, MED, and NED with three different β values.

fast execution time and proven effectiveness [9]. However, SED may present false negatives when, for example, 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 to SED's catalogued information entropy, even when anomalies are occurring. MED was designed to prevent false negative detections, refining the qualitative information entropy calculation by merging catalogued and monitored information into a merged list. As such, merged information entropy differs from catalogued information entropy whenever unregistered or missing elements occur in the monitored information. An example of a missing element is a VNF identifier not present in the monitored information of an SFC, but present in the catalogued information. Similarly, an unregistered element could be an additional port in the VNF monitored information that is not present in the catalogued information of that VNF.

For quantitative information, merging monitored and catalogued information is not mandatory due to their discrete nature, which makes it virtually impossible for the same entropy variations to appear in two different analyses. However, distinct entropy results may appear from every evaluation, making entropy-to-entropy comparison ineffective. NED was designed to overcome this issue, analyzing the monitored information entropy considering historical catalogued information entropies. For quantitative information,

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.

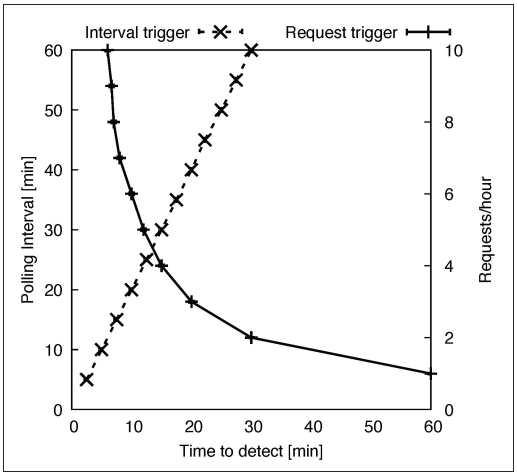

FIGURE 5. Detection time comparison of both triggers.

anomalies 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 β value, the more permissive the detector is, with β = 1 corresponding to no changes. The mechanisms receive as parameters the input data resulting from an adapted algorithm (described in detail earlier) and, in the case of NED, a β value. As such, it is imperative for datasets to be consistent because the anomaly detection accuracy relies on the integrity of the data used as input. The detectors' Python implementation is available at https://github.com/ComputerNetworks-UFRGS/nsm.

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 1 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 the VNF and SFC levels. Depending on the type of threat detected, it is possible to generate a cascade effect on other VNFs that are part of the same SFC. This cascade effect can also be detected since it generates 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 those 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 those 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 those stipulated, compromising the integrity security attribute.

Considering that every non-conforming pattern detected on VNF and SFC catalogued/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.

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:

1. The average number of SFCs (defined as 100, reflecting large-scale enterprise networks, with SFCs composed of 2 to 7 VNFs [14]).
2. 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]).
3. The threat event likelihood (defined as 60 percent based on enterprise reports [15]).
4. The legitimate event likelihood (also 60 percent). Both events' likelihoods follow a normal distribution.

The algorithm generates snapshots composed of two datasets:

1. Monitored data
2. Catalogued data

Each snapshot contains a timestamp, and each line of the datasets contains: SFC identifier, VNFs composing the SFC, 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, and 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 1. 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.

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

TABLE 1. Anomalies and threat characteristics.

With 95 percent on average, NED accuracy is slightly lower than MED, which results from scenarios where small variation in the monitored allocated bandwidth may not be detected, mainly at the start of NED execution, when few monitored values are available to compute the mean and standard deviation to define the monitored elements' entropy.

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 in the Values database, as mentioned earlier, and three values of β are considered:

- 1 (no change)
- 0.5 (half-size)
- 2 (double-size)

A receiver operating characteristic (ROC) curve is presented in Fig. 4, 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.

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

With 95 percent on average, NED accuracy is slightly lower than MED, which results from scenarios where small variation in the monitored allocated bandwidth may not be detected, mainly at the start of NED execution, when few monitored values are available to compute the mean and standard deviation to define the monitored elements' entropy. Thus, NED can present a high number of false positives until it has a certain number of samples. Afterward, 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. This means that greater entropy changes might be considered normal by NED although they might be anomalies. Still, using a smaller β may excessively limit analyzed samples, resulting in regular information potentially being considered anomalous.

Detection Time: Consider two NSM detection triggers: NSM internal interval and NFVO analysis request. 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, such as new VNF registration or configuration changes in existing VNFs. Such events' occurrences (regular or anomalous) varies from 1 to 10 per hour, following the 60 percent 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.

As the trigger interval increases, the anomaly detection delay increases linearly considering the *Interval* trigger. With the *Interval* trigger, anomalies are detected two times faster than using the polling interval on average. In turn, the time to detect anomalies with the *Request* trigger decreases logarithmically as the number of notifications per hour increases linearly. 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 could be detected faster without performing unnecessary analysis, which may occur when using short intervals for the *Interval* trigger.

The effectiveness of the *Interval* trigger relies on the interval configured. While short intervals may detect anomalies faster, this implies more NSM executions. In turn, while higher intervals imply fewer NSM executions, anomalies might take longer to detect. The *Request* trigger can be configured with different strategies, such as executing 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: acquire it, forward it to NSM, receive the anomaly detection results back, and evaluate the application (or not) of the suggested actions.

Conclusion and Future Research

This article discusses the advancements related to NFV environments' security, from its definition to recent proposals regarding NFV in emerging network environments. An NFV security module is presented to investigate anomaly detection effectiveness for NFV security. We analyze if threats related to security attributes could be properly detected using anomaly detection, which leads to the design, development, and evaluation of three distinct entropy-based anomaly detection mecha-

nisms. 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 percent among the entropy-based mechanisms designed. Also, two detection triggers are 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 *et al.*, "Network Function Virtualization: State-of-the-Art and Research Challenges," *IEEE Commun. Surveys & Tutorials*, vol. 18, no. 1, Mar. 2016, pp. 236–62.

[2] W. Yang and C. Fung, "A Survey on Security in Network Functions Virtualization," *IEEE NetSoft Conf. and Wksps.*, Seoul, South Korea, June 2016, pp. 15–19.

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

[4] M. Chiosi *et al.*, "Network Functions Virtualisation (NFV) — Use Cases," ETSI NFV ISG, Tech. Rep., 2013; https://www.etsi.org/deliver/etsi_gs/nfv/001_099/001/01.01.01_60/gs_nfv001v010101p.pdf, accessed Feb. 2022.

[5] L. Bondan *et al.*, "FENDE: Marketplace-Based Distribution, Execution, and Life Cycle Management of VNFs," *IEEE Commun. Mag.*, vol. 57, no. 1, Jan. 2019, pp. 13–19.

[6] T. Madi *et al.*, "NFV Security Survey in 5G Networks: A Three-Dimensional Threat Taxonomy," *Computer Networks*, vol. 197, 2021.

[7] M. Yousefi *et al.*, "A Reinforcement Learning Approach for Attack Graph Analysis," *IEEE Int'l. Conf. Trust, Security and Privacy in Computing and Commun./12th IEEE Int'l. Conf. Big Data Science and Engineering*, Aug. 2018, pp. 212–17.

[8] M. A. Ahmed and Y. A. Mohamed, "Enhancing Intrusion Detection Using Statistical Functions," *2018 Int'l. Conf. Computer, Control, Electrical, and Electronics Engineering*, Aug. 2018, pp. 1–6.

[9] L. Bondan *et al.*, "Anomaly Detection Framework for SFC Integrity in NFV Environments," *IEEE Conf. Network Softwarization*, Bologna, Italy, July 2017, pp. 1–5.

[10] L. Bondan and L. Z. Granville, *NFV Environments Security through Anomaly Detection*, Ph.D. dissertation, UFRGS, July 2019; https://lume.ufrgs.br/handle/10183/197460, accessed Oct. 2021.

[11] B. Briscoe *et al.*, "Network Functions Virtualisation (NFV) — NFV Security: Problem Statement," ETSI NFV ISG, Tech. Rep., 2014; https://www.etsi.org/deliver/etsi_gs/nfv-sec/001_099/001/01.01.01_60/gs_nfv-sec001v010101p.pdf accessed Feb. 2022.

[12] "ISO/IEC 27001: Information Technology — Security Techniques — Information Security Management Systems," ISO, Geneva, Switzerland, 2013.

[13] W. Rankothge *et al.*, "Data Modelling for the Evaluation of Virtualized Network Functions Resource Allocation Algorithms," *CoRR*, 2017; http://arxiv.org/abs/1702.00369, accessed Aug. 2021.

[14] J. Sherry *et al.*, "Making Middleboxes Someone Else's Problem: Network Processing as a Cloud Service," *ACM SIGCOMM Conf. Applications, Technologies, Architectures, and Protocols for Computer Commun.*, Helsinki, Finland, 2012, pp. 13–24.

[15] D. Anstee *et al.*, "Worldwide Infrastructure Security Report," Arbor Networks, White Paper, Tech. Rep., 2017; https://ir.netscout.com/investors/press-releases/press-release-details/2017/Arbor-Networks-12th-Annual-Worldwide-Infrastructure-Security-Report-Finds-Attacker-Innovation-and-IoT-Exploitation-Fuel-DDoS-Attack-Landscape/default.aspx, accessed Feb. 2022.

Biographies

Lucas Bondan is an R&D coordinator at the Brazilian National Research and Educational Network (RNP) and supervising professor at the University of Brasília. Among his research interests are NFV security, management and orchestration, and SFC.

Tim Wauters is a postdoctoral fellow at Ghent University. Management solutions for scalable multimedia delivery services and network and service architectures are some of his research interests.

Bruno Volckaert is a professor at Ghent University. His research interests include cloud computing advances, NFV/SFC (IETF standardization), 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, where he leads the network and service management research group. His research interests include scalable software architectures for network and service management, and the design and performance evaluation of novel QoE-aware multimedia delivery systems.

Lisandro Zambenedetti Granville is a full professor at the Federal University of Rio Grande do Sul. His research interests include management of network virtualization, intent-based networking, SDN, NFV, and network programmability.