A Security Monitoring Plane for Named Data Networking Deployment

Tan Nguyen, Hoang-Long Mai, Guillaume Doyen, Rémi Cogranne, Wissam Mallouli, Edgardo De Oca, Olivier Festor

► To cite this version:

Tan Nguyen, Hoang-Long Mai, Guillaume Doyen, Rémi Cogranne, Wissam Mallouli, et al.. A Security Monitoring Plane for Named Data Networking Deployment. IEEE Communications Magazine, 2018, 56 (11), pp.88-94. 10.1109/mcom.2018.1701135 . hal-02407673

HAL Id: hal-02407673

<https://hal.science/hal-02407673>

Submitted on 12 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Security Monitoring Plane for Named Data Networking Deployment

Tan Nguyen, Hoang-Long Mai, Guillaume Doyen, Rémi Cогranne,
Wissam Mallouli, Edgardo Montes de Oca and Olivier Festor.

Copyright ©2018 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org

Accepted version, final version available online on ieeexplore.ieee.org. DOI: 10.1109/MCOM.2018.1701135

Abstract—Named Data Networking (NDN) is the most mature proposal of the Information-Centric Networking (ICN) paradigm, a clean-slate approach for the Future Internet. Although NDN was designed to natively tackle security issues inherent to IP networks, it also introduces new security threats which may prevent its practical deployment by telco operators. Therefore designing and implementing a dedicated security monitoring plane is essential to enable such future deployment and in this paper, we present a set of contributions in this area. It first consists in featuring NDN significant attacks in a real operating context to evaluate their actual impact. Then, by analyzing the NDN Forwarding Daemon (NFD) data-plane pipelines, we present a monitoring plane design which captures the state of NDN nodes by instrumenting 18 metrics with dedicated probes. We then correlate these metrics with a Bayesian Network which allows the detection of potential abnormal behaviors. To validate our approach, we demonstrate the efficiency of our monitoring plane in the detection of Content Poisoning Attacks and Interest Flooding Attacks in a testbed carrying real traffic.

Index Terms—Named Data Networking, Bayesian Network, NDN monitoring, Network security, Anomaly detection.

I. INTRODUCTION

Despite its maturity and acknowledgment, Named Data Networking (NDN) [1] also introduces novel critical security issues [2] which may prevent it from replacing IP stacks in the near future. If one wants Internet Service Providers (ISPs) to adopt and deploy NDN in their operational infrastructures, a generic security monitoring plane is indispensable. Such plane

must be able to tackle not only existing security flaws as a whole but also threats that remain undiscovered to date.

In order to demonstrate the threat of NDN attacks in reality, we have implemented different attack scenarios and performed a comprehensive study of two significant NDN attacks – Interest Flooding Attack (IFA) and Content Poisoning Attack (CPA) – under real deployment conditions. Then, based on a thorough analysis of the NDN Forwarding Daemon (NFD) operations, we have proposed a set of 18 metrics which exhaustively feature the NDN data-plane. For each metric, we designed a statistical micro detector that raises alarms whenever a metric strays from its expected behavior. Our detector follows a design methodology which allows guaranteeing a prescribed false alarm rate. Then, the core of our monitoring plane leverages a Bayesian Network, a probabilistic graph-oriented approach that formulates causal relationships while handling uncertainty using probability theory. It allows us to correlate different metrics and use them to jointly infer anomalies from micro detectors’ alarms. To validate our monitoring plane, we have implemented our proposal into a real NDN probes with a monitoring operator. By using traffic data issued from a real deployment, we then demonstrate the performance of our micro detectors and the capability of our Bayesian Network to detect attacks under different scenarios.

The present paper partially summarizes our prior contributions [3]–[7] that focused on detecting specific attacks using dedicated metrics and an ad-hoc detection method. Furthermore, it also presents novel results from the implementation of IFA and CPA detectors in an industrial tool: Montimage Monitoring Tool (MMT)¹.

The paper is organized as follows. Section II highlights related works on NDN architecture and NDN security issues. Section III shows the result of a measurement campaign for featuring IFA and CPA. Next, Section IV presents the core of our contribution, including the NFD pipelines instrumentation, the micro detector design and the Bayesian Network that combines micro detectors’ alarms to infer anomalies. In Section V, by leveraging data collected on a real NDN testbed, we provide numerical results that demonstrate our approach

Copyright (c) 2018 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org

Accepted version, final version available online on ieeexplore.ieee.org. DOI: 10.1109/MCOM.2018.1701135

Tan Nguyen, Guillaume Doyen and Rémi Cогranne are with trans-disciplinary cyber-security team, ICD, UMR 6281 CNRS, Troyes University of Technology, Troyes, France. Guillaume Doyen is with the Autonomous Network Environment Team and Rémi Cогranne is also with the Lab. of System Modeling and Dependability.

Hoang-Long Mai, Wissam Mallouli and Edgardo Montes-De-Oca are with Montimage Research Labs, 39 rue Bobillot, 75013, Paris, France

Olivier Festor is with TELECOM Nancy, University of Lorraine, 54506 Vandoeuvre-les-Nancy, France, France

This work is partially co-funded by (1) the French National Research Agency (ANR), DOCTOR project, <ANR-14-CE28-0001>, started in 01/12/2014 and supported by the French Systematic cluster and (2) the CRCA and FEDER CyberSec Platform, <201304601>.

¹See: <http://www.montimage.com/products.html>. Accessed on November 10th, 2017.

relevance and performance. Finally, Section VI concludes the paper and presents our plans for future work.

II. RELATED WORKS

A. Information-Centric Networking and Named Data Networking

ICN is an emerging networking paradigm relying on content objects [8]. Instead of identifying IP addresses of hosts, ICN associates a content name to each content object. Among ICN proposals [8], NDN [1] is the most mature candidate to substitute for the current IP network. Communications in NDN are performed by *Interest* and *Data* packets. When one wants to get a specific content, he sends an *Interest*, and the content is returned in a *Data* packet. NDN also introduces three new router components. First, to improve the content delivery, the *Content Store* (CS) caches *Data* of recent requests. Secondly, the *Forwarding Information Base* (FIB) contains the routing information used to forward *Interests*. Finally, to forward the *Data* correctly, routers keep traces of all forwarded *Interests* in the *Pending Interest Table* (PIT). Whenever an *Interest* is forwarded, the associated incoming faces are recorded in a PIT entry. As such, when the corresponding *Data* arrives, it can be sent through the reverse path to the user. The router removes the corresponding PIT entry after forwarding the matching *Data*.

With more than 70 contributors involved in all NDN development projects, this solution is now mature enough for large deployments. In this effort, the NDN project provides NFD², a functional implementation of NDN forwarder along with libraries to develop NDN applications in multiple languages, thus encouraging researchers to study and experiment in NDN. Moreover, the global NDN testbed³ has been growing over the past few years and currently consists of 39 sites located all around the world.

B. Named Data Networking Security

As an approach for the Future Internet, NDN was designed to natively tackle security issues inherent in IP network. Nevertheless, its communication model and new router components expose the network to other attack types [2]. Among them, IFA and CPA are identified by the NDN community as the most significant ones⁴. IFA is a variation of the Denial of Service attack in NDN [9], which has drawn the attention of the NDN community from early stages [9]–[11]. Its principle is to send many *Interests* for nonexisting content. Such *Interests* cannot be satisfied and will occupy the PIT until expiration. When the PIT is full, the router cannot handle new incoming requests. Meanwhile, in CPA [12]–[14], a legitimate *Interest* is responded by a bad *Data* packet which can be inserted into the network by compromised routers or a collaboration between malicious providers and consumers. Bad *Data* has a valid content name, but its content is altered. Such an attack

leverages NDN caches to spread bad *Data* to as many users as possible and is likely to target popular content names to increase the attack impact.

III. FEATURING NDN ATTACKS IN A REAL DEPLOYMENT

As a first step toward the design and implementation of a security monitoring plane for NDN data-plane, we implemented the two most important attacks acknowledged by the NDN community, IFA and CPA, as described in Section II-B. This work allowed assessing the feasibility of these attacks, identifying reproducible attack scenarios, and characterizing their footprint on different components of the network (routers, users, and providers).

A. Interest Flooding Attack

While IFA was straightforward to implement at the early stage of NDN, its threat in real deployment remains unexplored, especially since the Negative ACKnowledgement (*NACK*) packet [15] was implemented in NFD. The *NACK* packet enables a router to notify downstream neighbors when it can neither satisfy nor forward an *Interest*, thus preventing unresolved *Interests* from lingering indefinitely in the PIT.

To answer this question, we implemented a basic setup consisting of a single NDN node and tried to overload the PIT in all possible ways [5]. These experimentations demonstrated that one could easily overload the PIT and eventually make NFD crash with the help of a malicious provider which delays *Data* packets for all *Interests* it receives. Besides increasing the number of clients, the attack impact can be exacerbated by leveraging long prefixes, regarding both the length and the number of levels in the naming hierarchy.

To characterize IFA in an NDN node, prior works [9]–[11] proposed using the “unsatisfied-*Interest*” (or unsatisfaction) ratio as the primary indicator for such an attack. This indicator is available by instrumenting the number of *Interests* sent and the number of *Data* received on node’s faces. As depicted in Figure 1, this novel attack pattern increases the unsatisfaction ratio significantly and, therefore, does not outdate current detection solutions.

B. Content Poisoning Attack

Following the same methodology, we have deployed an NDN topology, depicted in Figure 2, which is simple yet gathers all the basic components of a network operator. The purpose is to search for all means enabling CPA in reality. To the best of our knowledge, this is the first effort in this security area of NDN since the current state of the art either relies on strong assumptions (e.g., pre-polluted cache) or only considers CPA in simulated environments [13], [14]. We have eventually revealed attack scenarios in which the attacker competes with legitimate clients to prefetch nodes’ caches with poisoned contents by leveraging both the version number and the *Exclude* field of an *Interest*. We have highlighted the impact of these attack scenarios under the two NFD forwarding strategies, named *bestroute* and *multicast*. An alternative scenario, entitled *unsolicited*, exploits the fact that an NDN

²See: <http://named-data.net/doc/NFD/current>. Accessed on November 10th, 2017.

³See: <https://named-data.net/ndn-testbed>. Accessed on November 10th, 2017.

⁴See: <http://named-data.net/project/faq>. Accessed on November 10th, 2017.

Figure 1. IFA impact on unsatisfaction ratio

Figure 2. Use-case topology for Content Poisoning Attack.

node accepts *Data* even if it comes from a different face from the one the corresponding *Interest* has been forwarded to. This allows a bad provider to send *Data* packets unsolicitedly with a chance to match a pending *Interest* in the target NDN node.

Results show that while the IFA can mostly be featured using only the unsatisfaction ratio, it is more difficult to characterize CPA because its impact is not limited to a single metric and varies according to the location of the considered node in the topology. To entirely feature this effect, we performed a Principal Component Analysis (PCA), a well-known method for analyzing sets of measurements to find the most informative “axes,” which are a combination of metrics that best represents the observations. The result is depicted in Figure 3 using the two first components that account for 80.5 percents of data variance. The first component represents the CPA impact on bad *Data* injection, while the second one associates with the impact on the routers’ cache misses and the additional traffic to the provider. The figure shows

Figure 3. Projections of CPA measurements on axes of the two first principle components.

that the featuring of CPA heavily depends on the scenario, especially for the *unsolicited* scenario, because it has a very different footprint as compared to other scenarios. The PCA also shows that most of the metrics are impacted by the CPA: by generating additional traffic to put poisonous content in NDN caches, the CPA affects both the traffic behavior, the PIT usage and CS on each node. In other words, a single metric cannot feature the CPA impact entirely and thus is insufficient to enable an accurate CPA detection. This remark emphasizes the relevance of a generic monitoring plane that can raise alarms for any abnormal behavior, whether it is a currently known attack or a threat that remains undiscovered to date.

IV. DESIGN OF THE MONITORING PLANE

To design and implement a security monitoring plane for NDN networks, one first has to collect a comprehensive set of metrics in order to characterize a node’s operational state. Then, local detection algorithms must be designed to identify any deviant behaviors of a node. Finally, to deal with complex events, a correlation engine that is able to aggregate events from different node components and topology locations must also be addressed. In this section, we present our contributions in this area.

A. Featuring an NDN node Data Plane

Since we aim at designing a generic monitoring plane for the security of NDN data-plane, we propose to monitor a comprehensive set of metrics to provide as much information as possible on a node’s operational state. The proposed set results from a thorough analysis of packet processing described in Section II-A. We consider all relevant components inside an NDN node, including CS, PIT, and faces. The CS is instrumented to monitor the use of the cache through the number of content inserted, missed and hit on a periodic basis. Similarly, from the PIT, one can periodically monitor the number of *Interests* created, deleted, updated, as well as the average existing time of PIT entries before they are removed. Meanwhile, for each face, one can monitor the amount of incoming, outgoing and dropped *Interest*, *Data* and *NACK*

packets. We deliberately decided not to cover the FIB in our metrics list. Different from other components involved in our selected set, the FIB belongs to the control plane whose changes only occur due to static routing configurations or routing protocol announcements. Either way, its metrics are less likely to change as compared to those of other components and, thus, are less useful for featuring the node's behavior promptly. Moreover, we argue that the FIB malfunctions can be indirectly captured by other metrics. For instance, an increase in the number of dropped *Interests* and *Data* can mean that the FIB is not working correctly.

Current NFD implementation already proposes a management protocol⁵ which mostly features a node from a performance perspective with several basic metrics and the capability to configure an NDN node from a remote manager. This existing solution partially satisfies security requirements since it cannot provide fine-grained metrics (e.g., number of CS hits). Therefore, we have leveraged data from NFD logs to perform the node instrumentation. Altogether, this set of 18 metrics, is integrated into a local monitoring agent.

B. Micro Detector

The next step toward the design of a security monitoring plane, and indeed the most crucial part, consists in combining each measurement to get the relevant metric on the event that one wishes to detect and to design a method that decides when an abnormal event alarm should be triggered. The two main difficulties in the design of such a detector are the capability of (1) adapting with the natural dynamic aspect of network traffic and (2) modeling both normal and abnormal behavior of all the metrics.

The general idea of the method we used to design our micro detectors is based on constant Rate of False Alarm (RFA) tests that aim at (1) achieving a prescribed RFA and (2) maximizing the probability of detecting abnormal events under this false alarm constraint. We argue that controlling the RFA is crucial from an operational point of view to avoid numerous false alarms, and hence, to have a reliable detection of abnormal events. It is noteworthy that, in practice, maximizing abnormal event detection probability is hardly possible when numerous different events are considered. Therefore, most of our micro detectors are two-sided tests built upon a statistical model of legitimate traffic. For several of those metrics, an anomaly is expected to shift the metric in a specific direction. The unsatisfaction ratio, for instance, can only be increased under an IFA. In such case, a one-sided most-powerful detector may be designed. Examples of NDN attack micro detectors that are built on such statistical models and adaptable to dynamic aspect of traffic can be found in [3], [4].

C. Correlation Engine of Security Events

The proposed micro detectors, based on a single metric, can characterize only a specific aspect of an NDN node's status. Such limitation is highlighted in the case of CPA.

As indicated in Section III-B, CPA increases at the same time the PIT size, the number of packets on faces and the turnover in the CS due to the competition between bad and good content. Generally speaking, NDN attacks can differently impact nodes according to (1) their topological locations and (2) their internal components. Besides, it appears that alarms from a single micro detector cannot accurately detect and characterize an anomaly in an NDN network, thus making the combination of alarms from micro detectors essential.

To consider the causal relationships between micro detectors in our monitoring plane, we have proposed a Bayesian Network (BN) structure, depicted in Figure 4, whose nodes correspond to micro detectors associated with a given metric. BN is a probabilistic graph-oriented approach that formulates causal relationships while handling uncertainty using probability theory. Each node in its structure represents a random event/variable. A relation between two events is visualized by an edge between two corresponding nodes, starting from the parent ("cause" event) and pointing to the child ("affected" event). Nodes' relationships are quantified by their parameters, i.e., the probability of a child's value given values of its parents. Thus, given values of several nodes, one can infer the probability of a specific node's value. The reasons for our choice of BN are many. First, in BN, most of the events and their impact can be correlated on a small set of metrics. The repetition of the process for all metrics and all variables enables BN to use all of those relations to classify the final observed event, which is, in our case, the anomaly occurring in an NDN node. Secondly, as BN is a graphical model, the whole anomaly detection system can leverage a hierarchical approach with local metrics associated with local detectors, and a global detector acting as the root component for the network security. Finally, the metrics measured in computer networking are generally not entirely predictable. BN can efficiently deal with the underlying random nature of observed metrics using the Bayesian probabilistic approach.

In our BN, the Anomaly node represents the anomalies that can occur in an NDN network, and the directed edges are sketched based on NFD forwarding pipelines. A forwarding pipeline is a series of steps that operate on a packet or a PIT entry, triggered by a specific event as designed in the NFD node architecture specification⁶. For our security purpose, we group NFD pipelines in four main categories that are triggered by external factors: (1) incoming *Interest*; (2) unsatisfied *Interest*; (3) incoming *Data* and (4) incoming *NACK*. For instance, to partially motivate the edges provided in Figure 4, when an *Interest* arrives, NFD first checks if it violates reserved prefix and drops it, meaning that *Incoming Interest* impacts *Drop Drop Interest*. Afterward, if the *Interest* is duplicated with one that was already registered in the PIT, NFD sends a *NACK* message to notify the downstream. Otherwise, NFD inserts a new PIT entry or updates the corresponding one that already exists by canceling the unsatisfied timer. Hence, *Outgoing NACK*, *PIT Create*, and *PIT Update* are affected by *Incoming Interest*.

⁵See: <https://redmine.named-data.net/projects/nfd/wiki/Management>. Accessed on November 10th, 2017.

⁶See: NFD Developer's Guide - Revision 7. Accessed on November 10th, 2017

Figure 4. The proposed Bayesian Network.

V. DEPLOYMENT AND EXPERIMENTAL ASSESSMENT

The assessment of our methodology is presented in three steps. First, the experimental setup, implementing a real deployment of NDN environment with routers, consumers, and providers is introduced. Secondly, micro detectors are evaluated with IFA that exclusively impacts a single metric. Finally, the monitoring plane is evaluated with CPA. Although the monitoring plane was designed for a general purpose, only a handful of NDN attacks are discovered and, for most of them, their practical implementation in real NDN environment remains unexplored. This dramatically reduces the number of potential abnormal events that can be used to validate the monitoring plane.

A. Testbed and Scenarios

The experimental results are obtained within a virtualized environment that reproduces the topology presented in Figure 2. Such environment allows us to adjust the topology according to the need (e.g., adding users only requires spawning new containers). The topology consists of three routers: an edge router on the client side R1, a core router R2 and router R3 that represents the edge router and caching system on the legitimate provider side. Good clients and attacker clients connect to R1. In each router, an MMT probe is deployed to collect data for the 18 selected metrics. The MMT probe is implemented as an independent application that runs separately from NFD and extracts data from NFD logs. The probe consumes about 10% of memory and does not have a notable impact on the router. Then, the data is processed by micro detectors, which retrieve the node's parameters from a local database. Finally, the results of micro detectors are aggregated by the Bayesian Network to determine the abnormal event of the node.

Good clients' behavior is implemented by reproducing a realistic network traffic pattern which considers both the popularity of contents and the statistical properties of the requests over time. The attack is launched by dedicated bots.

The payload of the attack is controlled by the amount of traffic sent by bots as well as their number.

B. Micro Detector Validation: The Case of Interest Flooding Attack

To evaluate the efficiency of both the micro detectors and the monitoring plane, all experiments were carried out for 10 minutes. There is only normal traffic during the first 5 minutes. By contrast, during the last 5 minutes, the bots attack by sending malicious traffic. For IFA, those experiments were repeated 50 times to ensure the reliability of results.

As shown in prior works [3], [4], [9], [10] and stated in Section III.A, IFA almost exclusively affects the unsatisfaction ratio. Thus, this single metric can be used to detect it and evaluate the relevance of our micro detector design. The design approach for micro detectors relies on the modeling of each metric such that the micro detector raises the alarm only when the associated metric shift from the expected statistical model with a prescribed level of significance. For the unsatisfaction ratio used in IFA detection, after some processing [3], [4] and normalization, it can be modeled as a Gaussian (normal) distribution. Figure 5 shows that this model fits well with the statistical distribution of IFA's micro detector output. The figure also compares the micro detector's output when the IFA is occurring, with two different attack payloads. One can note that, while the micro detector's output is expected to increase under IFA, it mostly spreads in a much wider range. Such phenomenon emphasizes the complex reaction of real NDN deployment and also highlights the relevance of having general-purpose micro detectors that do not focus on detecting a specific event but instead alert when observations do not fit with the model. This is also illustrated in Figure 5 through the bounds of significance level, outside which value is considered as "abnormal" for the significance levels of 5% and 0.1%. With those thresholds, the empirical RFA is about 5.12% and 0.38%, respectively, while the IFA detection probability, i.e., true positive rate, is about 69.21% and 55.54%, respectively. This shows the excellent match between the expected and

Figure 5. Comparison between empirical and theoretical distribution of the transformed unsatisfaction ratio for legitimate and attack traffic.

practical RFA, as well as the very high level of attack detection which can be further increased by analyzing consecutive observations.

C. Monitoring Plane Assessment: The Case of Content Poisoning Attack

Among the three CPA scenarios mentioned in Section III-B, the *unsolicited* scenario can easily be fixed by patching NFD [6]. Indeed, since version 0.5.0 (October 4th, 2016), NFD has prevented unsolicited *Data* from being forwarded. Although the two other scenarios (i.e., *bestroute* and *multicast*) are more tedious to implement, they can hardly be circumvented as they make use of the NDN protocol and the NFD forwarding rules. Therefore, we leverage them to exhibit the relevance and efficiency of the proposed security monitoring plane.

Similarly to IFA, each CPA experiment lasts 10 minutes, and the attack is launched during the last 5 minutes. The bad provider in Figure 2 has a smaller delay than the legitimate provider but it does not take part in the default route and, therefore, it is associated with a higher cost. Legitimate users keep sending *Interests* until the good content is retrieved, using the *Exclude* field to avoid getting any previously received bad *Data*. Meanwhile, the attacker clients behave the other way round: they keep sending *Interests* excluding the good content so that the requests for the poisoned content names can be tricked to go to the bad provider. For a meaningful comparison, the average traffic amount from legitimate clients is set to 10 *Interests* per second. The attacker average traffic amount varies from 5 to 50 *Interests* per second, following a logarithmic scale. It is noteworthy that we also consider a scenario in which there is additional traffic from legitimate users [7]. Data collected from this scenario was added to the learning data set for the BN so that it can differentiate between the additional legitimate traffic and the malicious one.

The monitoring plane results for CPA detection are summarized in Table I. The table shows that the larger the attack payload, the better the detection. Even for a rather small attack payload, the monitoring plane can detect the CPA with rather high efficiency. It is noteworthy that those results correspond

Table I
EFFICIENCY OF MONITORING PLANE WITH RESPECT TO CPA DETECTION FOR TWO DIFFERENT SCENARIOS AND VARIOUS ATTACK PAYLOAD.

Scenario	Attack rate (# Interest/s)	5	10	20	50
CPA <i>bestroute</i>	% True Positive	95	95.33	97	98.33
	% False positive	<0.01	<0.01	<0.01	<0.01
CPA <i>multicast</i>	% True Positive	63.33	72.83	79.33	96.33
	% False positive	<0.01	<0.01	<0.01	<0.01

to the CPA detection done by correlating all metrics' measurement at a given instant, which is very powerful for real-time anomaly detection. Moreover, regardless of the attack payload and the scenario, the monitoring plane achieves a minimal amount of false positives (less than 0.01%) which is especially important in an operational context.

VI. CONCLUSION AND FUTURE WORKS

In this paper, we featured the two most significant NDN attacks, IFA and CPA, and proved their feasibility in an NDN real deployment. While a single metric can feature IFA, more insights of the NDN node status is needed to characterize CPA as well as unrevealed potential attacks in the future. This fact motivates the necessity of a security monitoring plane. To this aim, we proposed a comprehensive set of 18 NFD metrics based on a thorough analysis of NFD pipelines. For each metric, a micro detector was designed to capture any abnormal variation from the metric's normal behavior with a prescribed RFA. The relevance of the micro detector's design was evaluated through its performance against IFA in our testbed. For attacks that affect several aspects of NDN node status, like CPA, a correlation engine for monitored metrics based on a Bayesian Network is proposed to combine micro detector alarms in order to identify any abnormal security events in an NDN node. To validate these proposals, two CPA scenarios were considered in a real testbed that implements all of these contributions. Results demonstrated that our solution accurately detects these attacks with various rates.

Our future research directions will focus on: (1) addressing other attacks and larger topologies to assess the genericity of our approach; (2) continuing the implementation of our solutions into MMT probes to contribute to the secure deployment of NFD; and (3) distributing the BN to correlate different alerts from multiple nodes, thus enabling a potential trace-back mechanism.

REFERENCES

- [1] L. Zhang et al., "Named data networking," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, 2014, pp. 66–73.
- [2] R. Tourani et al., "Security, privacy, and access control in information-centric networking: A survey," *IEEE Communications Surveys & Tutorials*, 2017.
- [3] T. Nguyen, R. Cогranne, and G. Doyen, "An optimal statistical test for robust detection against interest flooding attacks in ccn," *IFIP/IEEE IM*, 2015, pp. 252–260.
- [4] T. N. Nguyen et al., "Detection of Interest flooding attacks in named data networking using hypothesis testing," *IEEE WIFS*, 2015, pp. 1–6.
- [5] H. L. Mai et al., "On the Readiness of NDN for a Secure Deployment: The Case of Pending Interest Table," *IFIP AIMS*, 2016, Springer, pp. 98–110.

- [6] T. Nguyen et al., "Content Poisoning in Named Data Networking: Comprehensive characterization of real deployment," *IFIP/IEEE IM*, 2017, pp. 72–80.
- [7] H. L. Mai et al., "Towards a Security Monitoring Plane for Named Data Networking and its Application against Content Poisoning Attack," full paper accepted, presented in *IFIP/IEEE NOMS*, on 25 April, 2018 in Taipei, Taiwan.
- [8] B. Ahlgren et al., "A survey of information-centric networking," *IEEE Communications Magazine*, vol. 50, no. 7, 2012, pp. 26–36.
- [9] A. Afanasyev et al., "Interest flooding attack and countermeasures in Named Data Networking," *IFIP Networking Conference*, IEEE, 2013, pp. 1–9.
- [10] A. Compagno et al., "Poseidon: Mitigating interest flooding DDoS attacks in named data networking," *IEEE LCN*, 2013, pp. 630–638.
- [11] K. Wang et al., "Detecting and mitigating interest flooding attacks in content-centric network," *Security and Communication Networks*, vol. 7, no. 4, 2014, pp. 685–699.
- [12] S. DiBenedetto and C. Papadopoulos, "Mitigating Poisoned Content with Forwarding Strategy," *INFOCOM WKSHPS*, 2016, pp. 164–169.
- [13] C. Ghali, G. Tsudik, and E. Uzun, "Needle in a Haystack: Mitigating Content Poisoning in Named-Data Networking," *NDSS Workshop SENT*, 2014.
- [14] D. Kim, S. Nam, J. Bi, and I. Yeom, "Efficient Content Verification in Named Data Networking," in *ACM ICN*, 2015, pp. 109–116.
- [15] C. Yi et al., "Adaptive forwarding in named data networking," *ACM SIGCOMM computer communication review*, vol. 42, no. 3, 2012, pp. 62–67.

Rémi Cогranne is an Associate Professor at Troyes University of Technology (UTT), France, since 2013. He has regularly been a visiting scholar at Binghamton University between 2014 and 2017. He received his PhD in Systems Safety and Optimization from UTT in 2011, since on, his research focus on hypothesis testing applied to image forensics, steganalysis, steganography and computer network anomaly detection which lead to more than 55 papers and 3 International patents.

Wissam Mallouli is currently a research and development project manager at Montimage, France. He received his PhD in computer science from Telecom and Management SudParis (France) in 2008. His topics of interest cover formal methods for monitoring of functional, performance and security aspects of networks and applications. He is working in several European and French research projects. He also participates to the program/organizing committees of numerous national and international conferences.

Tan Nguyen is a PhD student in Troyes University of Technology (UTT), France. His PhD is co-supervised by Dr. Rémi COGRANNE and Dr. Guillaume DOYEN. His research area focuses on security issues in Information Centric Networks and especially the NDN proposal. His PhD takes part of the DOCTOR project, started in December 2014 and funded by the French National Agency of Research (ANR).

Hoang-Long Mai received his master in Information Systems Security from Troyes University of Technology in 2016. He is currently a Ph.D. student in a CIFRE (Industrial Convention of Formation by Research) contract between Montimage, Troyes University of Technology and INRIA Lorraine. His Ph.D. topic focuses on the Autonomous Monitoring and Control of Virtualized Network Functions with an application to Named Data Networking.

Edgardo Montes de Oca graduated as a Computer and Electronics engineer 1985 from Paris XI, Orsay and DEA in Computers from Paris VI, Jussieu 1986. He was research engineer and leader in Euriware, and Alcatel's and Ericsson's Research centres. In 2004 he founded Montimage, a research oriented SME. His main interests include monitoring the security and performance of 4G/5G networks. He has published more than 30 papers, book chapters and patents.

Guillaume Doyen is an Associate Professor at Troyes University of Technology, France, since 2006. His current research focuses on the design of autonomous management solutions for the performance and security of content distribution and virtualized infrastructures. He published more than 50 papers in the network and service management community. He is a TPC member of high-venue conferences (IFIP/IEEE CNSM, NOMS, IM) and a co-chair of several events (AIMS, ManSDN/NFV).

Olivier Festor is Professor in Computer Science at the University of Lorraine and Director of the TELECOM Nancy, the Graduate Engineering School in Computer Science. Chair of IFIP TC6 WG 6.6 and IEEE COMSOC member, he is active for more than 25 years in the Network and Service Management scientific community. His research interest are in Network Security Monitoring and Configuration.