

SHIELD: A Novel NFV-based Cybersecurity Framework

G. Gardikis, K. Tzoulas, K. Tripolitis, A. Bartzas, S. Costicoglou
R&D Department
Space Hellas S.A.
Athens, Greece

B. Gastón, C. Fernández, C. Dávila
Open Big Data Technologies and Software Networks
research units.
Fundació privada I2CAT
Barcelona, Spain

L. Jacquin, H. Attak
Hewlett Packard Labs
Hewlett Packard Enterprise
Bristol, United Kingdom

D. Katsianis, I. Neokosmidis
INCITES Consulting SARL
Strassen, Luxembourg

T. Batista, R. Preto
UBIWHERE Lda.
Aveiro, Portugal

Antonio Lioy
Politecnico di Torino
Torino, Italy

A. Litke, N. Papadakis, D. Papadopoulos
R&D Department
Infili Technologies PC
Athens, Greece

A. Pastor, J. Nuñez
Telefonica I+D
Madrid, Spain

N. Davri, G. Xylouris, M. Kafetzakis
ORION Innovations P.C.
Athens, Greece

M. Terranova, C. Giustozzi
Agenzia per l'Italia Digitale
Rome, Italy

E. Trouva, Y. Angelopoulos, A. Kourtis
Institute of Informatics and Telecommunications
NCSR "Demokritos"
Aghia Paraskevi, Greece

Abstract— SHIELD is an EU-funded project, targeting at the design and development of a novel cybersecurity framework, which offers security-as-a-Service in an evolved telco environment. The SHIELD framework leverages NFV (Network Functions Virtualization) and SDN (Software-Defined Networking) for virtualization and dynamic placement of virtualised security appliances in the network (virtual Network Security Functions – vNSFs), Big Data analytics for real-time incident detection and mitigation, as well as attestation techniques for securing both the infrastructure and the services. This paper discusses key use cases and requirements for the SHIELD framework and presents a high-level architectural approach.

Keywords—*cybersecurity; NFV; big data analytics; infrastructure and service attestation*

I. INTRODUCTION

Cybercrime techniques continuously evolve to target victims and to subvert information technology. It is expected that mobility and heterogeneity of devices, as well as Big Data environments, will be two of the main targets for cybercrime in the years to come. The 2013 Norton Report [1] estimated at 13 billion dollars the economic losses due to consumer cybercrime just for Europe. Moreover, Ponemon study also points out that the tendency of these economic losses is increasing [2].

The priorities of the EC Digital Agenda state that protection against online accidents and crime has become central to consumer confidence and the online economy [3]. The success of attacks carries too many negative consequences for the victims, where most of these concerns regard the loss of sensitive data and intellectual property, opportunity costs (including service and employment disruptions), the damage to the brand image and company reputation, penalties and

contractual compensations to customers of commercial networks after service disruptions. Other consequences involve the need for costly countermeasures and insurance, the need for mitigation strategies and recovery from cyber-attacks, loss and/or distortion of trade and competitiveness and loss of work carried out [4]. All these concerns call for an effective strategy against cyber-attacks that accurately transforms shared knowledge into actionable information while maintaining a global view of the network.

SHIELD (Securing against intruders and other threats through an NFV-enabled environment) [5] is a recently launched EU-funded project with the ambition to address the above mentioned challenges by designing and implementing an integrated framework for next-generation security-as-a-service (SecaaS) offerings. The next sections of the paper describe the SHIELD concept, use cases, technical and business requirements as well as the high-level architecture of the SHIELD framework.

II. THE SHIELD CONCEPT

The SHIELD framework combines Network Functions Virtualisation (NFV), Security-as-a-Service (SecaaS), Big Data Analytics and Trusted Computing (TC), in order to provide an extensible, adaptable, fast, low-cost and trustworthy cybersecurity solution. It aims at delivering IT security as an integrated service of virtual network infrastructures, which can be tailored for Internet Service Providers (ISPs) and enterprise customers - including SMEs - in equal terms. Virtualised Network Security Functions (vNSF) provide software instantiations of security appliances that can be dynamically deployed into a network infrastructure. In line with the NFV concept and going beyond traditional SecaaS offers, vNSFs can be distributed within the network infrastructure close to the user/customer. This allows to radically optimize resource allocation, minimize costs and reduce incident response time.

Furthermore, SHIELD decouples security policies enforcement from their configuration; the security controller can dynamically – and remotely – reconfigured the vNSFs based on recommendation issued by big-data analytics. Leveraging the open approach of SHIELD, the core security components – vNSFs, security analytics and recommendations – are dynamically modifiable: security analyst can thus quickly develop and deploy new components for the SHIELD framework to adapt the platform to new type of attacks.

By separating the control plane from the enforcement (security) and data plane (network), a security gap may arise between an operator configuring the platform and the component implementing the configuration. SHIELD addresses this security issue by leveraging TC methods and technologies: the virtualisation software stack and the vNSFs running on it are measured and attested against their expected state. Similarly, the Software-Defined Network (SDN) used to steer the users' packets through their vNSFs is also attested to ensure the correct vNSFs chain is applied for each user. The application of attestation mechanisms to NFV is a topic which is gaining increasing interest by the community [6].

In current vNSF offerings, such as virtualised edge appliances, which operate independently [7], SHIELD

envisages that data and logs from vNSFs are aggregated and fed into an information-driven Intrusion Detection and Prevention System (IDPS) platform called Data Analysis and Remediation Engine (DARE), featuring analytical components capable of predicting specific vulnerabilities and attacks. The DARE relies on continuous monitoring of the network traffic, using monitoring vNSFs, and translates their observations into adversarial options, behaviours and intents. Centralising events and logs from multiple vNSFs, the DARE maintains the “big picture” of the network infrastructure status; thus it can infer events which cannot be detected by the individual vNSFs. Security modules within the DARE are responsible for analysing the monitoring data; the DARE supports multiple security modules, which can use different analytics methods to detect attacks. When a security module detects an attack, it provides possible recommendations for remediation and mitigation, such as updating a firewall vNSF configuration to block the attacks. Automatic mitigation in a highly diverse environment is among the features which differentiate the DARE from commercial SIEM platforms [8]. Monitoring information, events and notifications, overall security status and recommended actions become available through the security dashboard.

III. SHIELD USE CASES

The SHIELD framework brings together all actors in the security value chain (ISPs, enterprises, end users, cybersecurity agencies, security vendors) into a single ecosystem and facilitates the interactions between them, enabling new business models. Three main Use Cases (UCs) are foreseen for SHIELD.

A. Use Case 1: An ISP using SHIELD to secure their own infrastructure

In order to protect their own network infrastructure, ISPs have to deploy specific hardware which is very expensive since this hardware has to be maintained by very specialized operators. Furthermore, the operators may need to invest time troubleshooting the attack first. The virtualization offered by SHIELD in this use case aims to dramatically reduce both costs and response times by replacing specific hardware for vNSFs, as well as providing a central interface (dashboard) to understand the implications of the gathered information and analysis, and then act in the network.

B. Use Case 2: An ISP leveraging SHIELD to provide advanced SecaaS services to customers

SHIELD provides an ideal foundation for building enhanced SecaaS services, far beyond current offers. Using the SecaaS paradigm, the complexity of the security analysis can be hidden from the client (either a big company or an SME) who can be freed from the need to acquire, deploy, manage and upgrade specialised equipment. In this UC, the ISP can insert new security-oriented functionalities directly into the user local network, through its provided gateway or in its own network infrastructure.

C. Use Case 3: Contributing to national, European and global security

The dashboard, available to authorised actors, accepts ad-hoc requests regarding threat models or acquired threat intelligence. This data can be retrieved by, for instance, public cybersecurity agencies. The secure SHIELD framework offers, in this manner, a way of sharing threat information with third-parties who wish to synchronise information and research on measures to be taken for recent attacks, suffered by others. Currently, if a cybersecurity agency wants to retrieve statistical information about a network, it has to agree with the SP and deploy specific hardware on the infrastructure. This is a very costly procedure in both time and money, which makes it prohibitive for the current market situation. Attacks are constantly evolving and require a fast reactive and flexible solution. Using SHIELD instead, cybersecurity agencies can establish agreements with the SP and deploy vNSFs quickly and without extra cost in the infrastructure. Moreover, the analysed data is accessible from the dashboard because its processing is done in the DARE.

IV. REQUIREMENTS AND USE CASE PRIORITISATION

Following the use case definition, the next step is the identification of the high-level system requirements, which would drive the design task. For the gathering of the requirements, three sources were used:

- The three identified use cases (previously described in Section III)
- User stories, drafted from various stakeholders inside the SHIELD consortium expressing desired functionalities/interactions with users
- An online survey, aimed at prioritizing the use cases and collecting additional requirements.

The online survey was addressed at targeted persons, both within and outside the SHIELD consortium, who are professionally engaged with information security tasks. It was divided in three parts: profiling of the experts, criteria comparison and organizational aspects. The criteria comparison part used the Analytic Hierarchy Process (AHP) methodology in order to prioritise the three use cases based on several criteria.

AHP [9][10] is a structured technique for dealing with complex decisions based on a rational and comprehensive framework for decomposing an unstructured complex problem into a multi-level hierarchy of interrelated criteria, sub-criteria and decision alternatives. By incorporating judgments on qualitative and quantitative criteria, AHP manages to quantify decision makers' preferences. In the SHIELD survey, for each Use Case, the following categories of criteria were taken into account:

- Relevance of the use cases – Social and economic impact of the use cases (for the organisation, the EU market, and the EU society)

- Threats and vulnerabilities – Targeted threats or vulnerabilities addressed by the solution. (e.g. denial of service, data leakage, identity theft etc.)
- Security solution aspects – Aspects that cybersecurity solutions must address (e.g. cost, easiness to use, etc.)

The web-platform was implemented using Lime Survey [11], an open source tool for web surveys. (Fig.1). Overall, responses from 26 security experts (from both the academic and commercial sector) were recorded and analysed.

The result of the analysis of the responses shows that, among the use cases described in the previous section, UC2 (“An ISP leveraging SHIELD to provide advanced SecaaS services to customers”) is of higher value. It is preferred by half of the responders (mainly Businesses), followed at a distance by UC1 and UC3. The criteria identified as of high importance for the SHIELD platform are protection against data leakage and Identity theft, as well as compliance with organizational needs and policies. On the contrary, the less important identified aspects, among the listed ones, seem to be operational transparency and ease of use. Finally, the main results of the responses regarding the organizational aspects part of the survey show a good predisposition to deploy security services in a cloud environment (around 93% positive responses). The responders indicated the flexibility and the cost-efficiency as positive factors, but they also showed their concern regarding the service security.

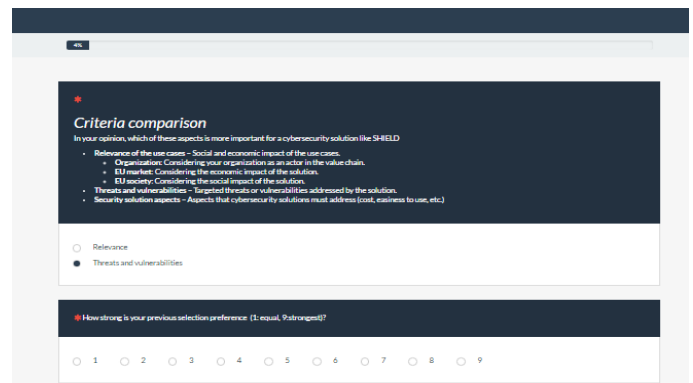


Fig. 1. Online survey for prioritising the use cases and collecting system requirements for SHIELD

The requirements elicited from the above mentioned sources are divided in i) general platform requirements and ii) vNSFs and analytics required.

In the first category, general functional requirements of the SHIELD platform are included, such as: vNSF deployment and lifecycle management, data monitoring, analytics and visualisation. Non-functional requirements for the SHIELD platform are also identified, concerning responsiveness, availability, and scalability.

In the second category, the functionalities needed by the vNSFs and the DARE's security modules are included. Based on the survey results, the most popular functionalities include: blocking the access to malware and malicious websites, Layer

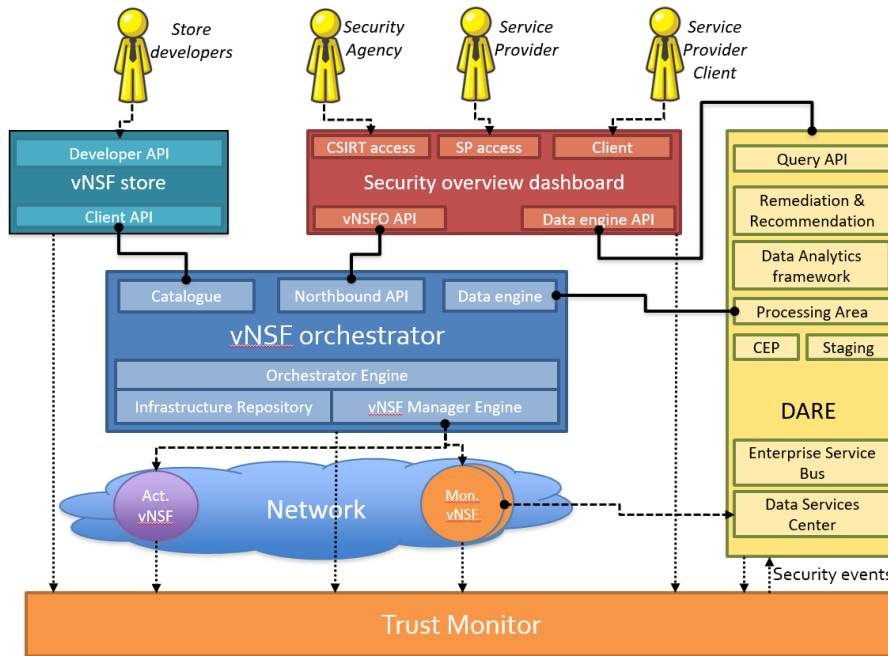


Fig. 2. SHIELD high-level system architecture

4 traffic filtering, spam protection, Distributed Denial of Service (DDoS) protection as well as Intrusion Detection System/Intrusion Prevention System (IDPS) functionalities.

V. THE SHIELD SYSTEM ARCHITECTURE

Based on the use cases and requirements highlighted in the previous sections, it is possible to draft an initial high-level architecture for the SHIELD framework. The architecture is articulated around different components, illustrated in Fig. 2 and described more deeply in this section.

A. Network infrastructure

The network infrastructure provides a trusted environment for supporting the execution of vNSFs. For these purposes, the infrastructure supports attestation and interacts with the Trust Monitor, which attests of the integrity of each network component.

Additionally, in order to be able to host the vNSFs, the network infrastructure also implements a Network Functions Virtualisation Infrastructure (NFVI) environment. The NFVI, according to the ETSI NFV specifications [12]-[14], includes the physical and virtual nodes (commodity servers, VMs, storage systems, switches, routers etc.) on which the services are deployed.

B. Virtual Network Security Functions (vNSFs)

vNSFs are software instantiations of security appliances that are dynamically deployed into the network infrastructure. There are two main types of vNSFs operating on the network. The first one is the monitoring vNSFs, devoted to gather information about the network traffic, and generate events sent to the DARE. The second type is the acting vNSFs, which

prevent attacks or mitigate vulnerabilities and threats. The proper vNSF solution is chosen depending on the kind of threat to defend against, and the associated security modules in the DARE.

In terms of vNSF architecture, the main differentiating factor in SHIELD from other NFV frameworks is the addition of the attestation capability (via the Trust Monitor) to the platform and the use of security analytics and a security controller.

C. vNSF Orchestrator

The vNSF orchestrator, or vNSFO, is responsible for managing the lifecycle of Network Services (NS), which are composed by one or more vNSFs. Among others, this allows to deploy (instantiate and place) vNSFs in specific points of the network infrastructure.

To that end, the vNSFO interacts with each of the other modules to obtain data on the vNSFs, to receive deployment requests or to convey information of specific vNSFs in order to enable analysis processes. The orchestrator also communicates with the infrastructure manager to deploy any requested vNSF or entire NS.

The orchestrator features some prominent sub-systems:

- The vNSF Manager handles the lifecycle of the vNSFs (provisioning and instantiation, configuration and update of parameters, scaling, termination etc.)
- The Catalogue sub-system, which includes information for both on-boarded vNSFs (vNSF descriptor, images) and NSs (NS descriptor, virtual link descriptor, vNSF forwarding graph).

- Two different Repositories containing the running instances for both vNSFs and NSs; and a relation of the NFVI (NFV Infrastructure) resources, properly modelled to use by the platform.
- The vNSF Monitoring module, which monitors the running vNSFs.

The vNSFO used in SHIELD will be based on the TeNOR Orchestrator, as developed from the FP7 T-NOVA project [15]. The vNSFO will follow the specification of ETSI NFV MANO [16].

D. vNSF store

The vNSF store acts as a nexus between the vNSFO and third-party vNSF providers/developers, who can register and manage vNSFs in order to make them available through the SHIELD platform. The following vNSF data are provided to and handled by the store:

- Service descriptors that contain developer information or versioning information (metadata), but also technical details concerning deployment requirements.
- Software images that contain the actual virtual appliances to be instantiated.
- Security descriptors, which contain information required to validate the integrity of themselves as well as the remaining files that comprise the service at all the critical procedures, such as on boarding, deployment and runtime.

E. Trust Monitor

The Trust Monitor is the component in charge of monitoring the trust of the SHIELD infrastructure. This is achieved by a combination of authentication and integrity: each node joining the infrastructure must be properly authenticated and provide also a proof of the integrity of its software stack, by leveraging TC mechanisms.

Integrity is also checked periodically to detect compromised software and if so, the vNSF Orchestrator is timely informed to take appropriate action (typically to quickly isolate the compromised node and reconfigure the infrastructure to maintain its expected functionality).

Integrity monitoring is based on the Trusted Computing paradigm and its Remote Attestation workflow. Each node is equipped with a TPM chip to provide a hardware root of trust. Additionally, suitable software is installed to measure all the relevant actions (from the boot phase up to the applications) and to report them in a secure and trusted way.

F. Data Analysis and Remediation Engine

The Data Analysis and Remediation Engine (DARE) is an information-driven IDPS platform that stores and analyses heterogeneous network information, previously collected via monitoring vNSFs. It features cognitive and analytical components capable of predicting specific vulnerabilities and attacks. The processing and analysis of large amounts of data is carried out by using Big Data, data analytics and machine

learning techniques. By processing data and logs from vNSFs deployed at specific strategic locations of the network, the DARE components provide feedback to cybersecurity data topologies and, in the case malicious activity is detected, they implement remediation activities, either by recommending actions by means of a dashboard and accessible API, or by (optionally) triggering task-specific countermeasures. The DARE platform provides flexible support for both new security capabilities and reconfiguration of existing security controls and allows extensions with multiple data analytics engines by providing a clear API to work with the collected data.

The DARE consists of three main components, the data collection and preparation module, the Data Analytics Engine and the Remediation Engine.

The *Data Collection and Preparation* module is responsible of the ingestion of the selected datasets and their preparation for further processing. Data to be collected include: flow information, DNS and proxy information, vNSFs logs and events and (in some cases) generic vNSFs monitoring metrics and status.

The *Data Analytics Engine* leverages different analytics modules (while opening the platform for the inclusion of others in the future) that use a wide range of complementary detection techniques along with open source frameworks and solutions. The Data Analytics engine is able to produce packet and flow analytics by using scalable machine-learning techniques. To this end, it involves the latest distributed computing technologies (Apache Spot, Spark, Storm, HDFS, Kafka) to allow for streaming processing of large amounts of data, scalability and load balancing, open data models and concurrent running of multiple machine-learning applications on a single, shared, enriched data set. The threat detection procedure of the cognitive module is based on the Apache Spot [17] framework. Specifically, the ingested data is available for searching, for use by machine learning, to be transferred to law enforcement, or as an input to other systems. Subsequently, the system uses a combination of machine learning tools to run scalable machine learning algorithms, not only as a filter for separating bad traffic from benign one, but also as a way to characterize the unique behaviour of network traffic. Finally, and in addition to machine learning, several processes of context enrichment, noise filtering, whitelisting, and heuristics are applied to network data, in order to present the most likely patterns that may comprise security threats.

Finally, the *Remediation engine* uses the analysis from the data analytics modules and is fed with alerts and contextual information to determine a mitigation plan for the existing threats. It performs in real-time or near-real-time, using open-source technologies (e.g. Storm). The Remediation Engine's main goal is to incorporate a combination of recommendations and alerts, which provide relevant threat details using the dashboard and the direct application of countermeasure activities by triggering specific vNSFs via the vNSFO (e.g. block/redirection of network flows).

G. Security dashboard and controller

The SHIELD platform provides an intuitive and appealing graphical user interface allowing SHIELD authenticated and

authorized users to access SHIELD's security dashboard. From this dashboard, operators have access to monitoring information showing an overview of the security status. The dashboard also allows operators as well as tenants to take actions and react to any detected vulnerability. Billing features is also present in the security dashboard allowing providers to measure and charge operations made by clients (for instance, the acquisition/instantiation of a new vNSF).

VI. USE CASE IMPLEMENTATION

With the designed architecture and the identified subsystems, it is possible to describe the sequences and interactions needed for the implementation of the defined Use Cases of Section III. For example, Use Case 2, can be implemented via the following steps (as shown in Fig. 3):

1. *Develop*: The vNSFs in this case can be developed by either the ISP or by a third party (vNSF developer). Once the vNSF are developed they can be deployed to the vNSF store.

2. *Study security services offered*: The client studies, using the dashboard, the security services offered by the SP.

3. *Select security services*: The client selects the desired security services.

4. *Acquire & deploy*: The client deploys, using the dashboard, the selected services, which may consist in one or more vNSFs that are located in the edge of its infrastructure.

5. *Gather monitoring information*: The deployed vNSF sends monitoring information to the DARE which acquires and validates the incoming data and then stores and processes them.

6. *Perform analytics*: The DARE processes the data according to the needs of the security services deployed by the client.

7. *Analyse & Recommendations*: The dashboard provides to the client monitoring data from the deployed security services and proposes mitigation actions.

8. *Acquire & deploy*: Depending on the recommendations and the security requirements, more vNSFs can be deployed in the infrastructure to protect the client.

VII. CONCLUSIONS

This paper presents an overview of a novel integrated framework for cybersecurity, being developed by the SHIELD project. This includes a definition of the prominent use cases, the identification of requirements, as well as a high-level architecture design of the SHIELD framework. The requirements collected via the online survey contributed to design a technical solution, which is well aligned to both the market needs and the recent trends in NFV architectures and big data analytics. In addition, the proposed architecture is compliant with the current technical approach as well as the terminology of ETSI ISG NFV.

Finally SHIELD's open design allows the framework to support new vNSFs and security analytics and remediation, which will protect against currently unknown attacks.

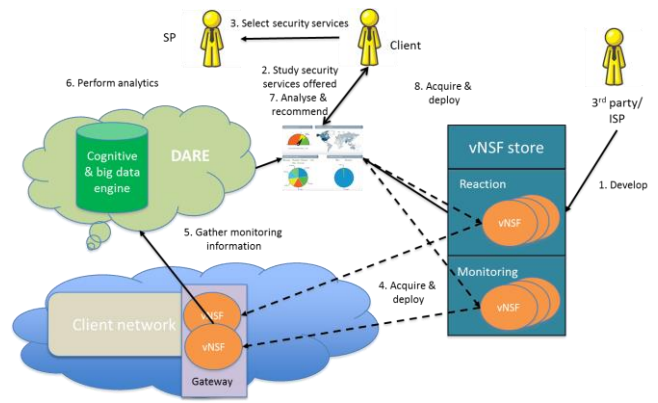


Fig. 3. Realisation of Use Case 2 via the proposed architecture

ACKNOWLEDGMENT

The work described in this article has received funding by the European Union Horizon 2020 research and innovation programme, under Grant Agreement no. 700199.

REFERENCES

- [1] P. Paganini, "2013 Norton Report, the impact of cybercrime according Symantec", Security Affairs, Oct. 2013, available online at <http://securityaffairs.co/wordpress/18475/cyber-crime/2013-norton-report.html>
- [2] Ponemon Institute, "2014 Global Report on the Cost of Cyber Crime", HP security, available online at <http://www8.hp.com/us/en/software-solutions/ponemon-cyber-security-report>
- [3] European Commission, "Digital Agenda for Europe", Nov. 2014, available online at http://europa.eu/pol/pdf/flipbook/en/digital_agenda_en.pdf
- [4] P. Paganini, "2013 - The Impact of Cybercrime", InfoSec Institute, Nov. 2013, available online at <http://resources.infosecinstitute.com/2013-impact-cybercrime>
- [5] EU SHIELD project, (Securing against intruders and other threats through an NFV-enabled environment) <https://www.shield-h2020.eu/>
- [6] S. Ravidas, S. Lal, I. Oliver and L. Hippelainen, "Incorporating trust in NFV: Addressing the challenges," 2017 20th Conference on Innovations in Clouds, Internet and Networks (ICIN), Paris, 2017, pp. 87-91. doi: 10.1109/ICIN.2017.7899394
- [7] RAD Comprehensive vCPE Toolbox, Comprehensive vCPE Toolbox
- [8] IBM QRadar SIEM, <http://www-03.ibm.com/software/products/en/qradar-siem>
- [9] T. L. Saaty, "A scaling method for priorities in hierarchical structures," Journal of Mathematical Psychology, vol. 15, pp. 234-281, 1977.
- [10] A. M. A. Bahurmoz, "The analytic hierarchy process at DarAl-Hekma, Saudi Arabia," Interfaces, vol. 33, pp. 70-78, 2003.
- [11] LimeSurvey, <https://www.limesurvey.org/>
- [12] ETSI NFV ISG. ETSI GS NFV 002 v1.1.1 Network Functions Virtualisation (NFV); Architectural Framework. s.l.: ETSI, 2013.
- [13] ETSI GS NFV 003 v1.1.1 Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV. s.l.: ETSI, 2013.
- [14] ETSI GS NFV-PER 002 V1.1.1 Network Functions Virtualisation; Proof of Concepts; Framework. s.l.: ETSI, 2013.
- [15] EU T-NOVA project, (Network Functions as-a-Service over Virtualised Infrastructures), <http://www.t-nova.eu>
- [16] ETSI GS NFV-MAN 001 (2014-12), Network Functions Virtualisation (NFV); Management and Orchestration
- [17] Apache Spot project, <https://spot.apache.org>