



City Research Online

City St George's, University of London

Citation: Petroulakis, N. E., Fysarakis, K., Askoxylakis, I. G. & Spanoudakis, G. (2018). Reactive Security for SDN/NFV-enabled Industrial Networks leveraging Service Function Chaining. Transactions on Emerging Telecommunications Technologies, 29(7), e3269. doi: 10.1002/ett.3269

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/18581/>

Link to published version: <https://doi.org/10.1002/ett.3269>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Reactive Security for SDN/NFV-enabled Industrial Networks Leveraging Service Function Chaining

Nikolaos E. Petroulakis^{1,2}, Konstantinos Fysarakis¹, Ioannis Askoxylakis¹ and George Spanoudakis²

¹Foundation for Research and Technology-Hellas, Greece ²City, University of London, United Kingdom

ABSTRACT

The innovative application of 5G core technologies, namely Software Defined Networking (SDN) and Network Function Virtualization (NFV), can help reduce capital and operational expenditures in industrial networks. Nevertheless, SDN expands the attack surface of the communication infrastructure, thus necessitating the introduction of additional security mechanisms. These major changes could not leave the industrial environment unaffected, with smart industrial deployments gradually becoming a reality; a trend that is often referred to as the 4th industrial revolution or Industry 4.0. A wind park is a good example of an industrial application relying on a network with strict performance, security, and reliability requirements, and was chosen as a representative example of industrial systems. This work highlights the benefit of leveraging the flexibility of SDN/NFV-enabled networks to deploy enhanced, reactive security mechanisms for the protection of the industrial network, via the use of Service Function Chaining. Moreover, the implementation of a proof-of-concept reactive security framework for an industrial-grade wind park network is presented, along with a performance evaluation of the proposed approach. The framework is equipped with SDN and Supervisory Control and Data Acquisition (SCADA) honeypots, modelled on and deployable to the wind park, allowing continuous monitoring of the industrial network and detailed analysis of potential attacks, thus isolating attackers and enabling the assessment of their level of sophistication. Moreover, the applicability of the proposed solutions is assessed in the context of the specific industrial application, based on the analysis of the network characteristics and requirements of an actual, operating wind park.

1. INTRODUCTION

With anticipated exponential growth of connected devices, future networks require an open-solutions architecture, facilitated by standards and a strong ecosystem. Such devices need a simple interface to the connected network to request the kind of communication service characterized by guarantees about bandwidth, delay, jitter, packet loss or redundancy. In response, the network should grant the requested network resources automatically and program the intermediate networking devices based on device profile and privileges. A similar requirement also comes from business applications where application itself asks for particular network resources based on its needs. Software Defined Networking (SDN) and Network Function Virtualization (NFV), important parts of 5G networking provide promising combination leading to programmable connectivity, rapid service provisioning and service chaining and can thus help lower capital and operational expenditure costs (CAPEX/OPEX) in the control network infrastructure. Nevertheless, SDN and NFV expand

the attack surface of the communication infrastructure, necessitating the introduction of additional security mechanisms. Industrial networks typically come with strict performance, security, and reliability requirements. SDN in the *Industry 4.0* arrives as a new concept to promote the computerization of the manufacturing part of the network that can help in communicating essential technologies such as the Internet of Things (IoT), communication machine-to-machine (M2M) and Cyber-Physical Systems (CPS) [1]. Furthermore, by appropriately leveraging the flexibility of SDN/NFV-enabled networks in the context of the adopted security mechanisms, industrial infrastructures can not only match but also improve their security posture compared to the existing, traditional networking environments [2].

This paper showcases a representative use case of an industrial network by considering an industrial control network for wind park operations. The wind park control network has been chosen as a key industrial application as wind energy has now established itself as a mainstay of sustainable energy generation. Nevertheless, the flexibility

of SDN networks means they can also help provide better security for industrial networks. Due to the controller's global view of the network and the ability to reprogram the data plane at real-time, SDN allows not only to revisit old security concepts (e.g. firewalls) but introduce new techniques as well (e.g. steering suspicious traffic to Supervisory Control and Data Acquisition (SCADA) Honeypots, adopting moving target defense and other reactive techniques). The deployment of these enhanced security concepts is in line with the enhanced protection requirements of critical infrastructures, given that the old paradigm of perimeter defenses and trusted internal networks is obsolete, as recent attacks have demonstrated [3]. Thus, enhanced security services are more than "good practice", but a requirement, as evidenced, for example, by the recent update to North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection standards, such as the measures detailed in the latest versions of CIP-007 (i.e. CIP-007-6 [4]), which dictate continuous network monitoring and deployment of network defenses to detect/block malicious activity within the Utilities' perimeter.

Service Function Chaining (SFC) provides the ability to define an ordered list of network services [5]. The concept of SFC has shown promising results providing the ability to define an ordered list of a network services to create a service chain. These services are then "stitched" together in the network to create a service chain, allowing us to route unknown/suspicious traffic via the Intrusion Detection and Deep Packet Inspection service functions, to classify it (as either legitimate or malicious), allowing us to forward it to the wind park or the honeypot, accordingly. With this mechanism, malicious traffic can be isolated in the honeypot, allowing us to track the attacker, identify her purpose and keep her occupied. Motivated by the above, we present a Reactive Security Framework for next generation 5G (and SDN/NFV in specific) -enabled industrial networks leveraged by SFC. More specifically, considering the energy production critical infrastructures, the framework features enhanced security functions, such as SCADA honeypots, are modelled based upon an operational wind park and ready to be deployed in one. The presented framework allows the continuous monitoring of the wind park industrial network, with provisions to reduce the impact of the security functions on the network's performance and to alleviate the burden of deploying and managing the security services themselves.

Moreover, the framework's Honeynet (consisting of both an active SCADA-specific honeypot and a passive honeypot) facilitates the detailed analysis of potential attacks, isolating attackers and enabling the assessment of their level of sophistication (e.g. from script kiddies to state actors). Building upon the concept presented in [6], this work highlights various use cases where the proposed mechanisms would be useful in the context of actual wind park deployments and associated business requirements. Moreover, a full implementation of the framework is presented, along with a performance evaluation, on a realistic testbed featuring various services and operational security service functions. The results of this evaluation are assessed in the context of an actual industrial network, highlighting the most viable options in the context of a real industrial application. Said assessment is based on a trace analysis conducted in an operating wind park's network (in

Brandø, Denmark); a representative use case of industrial networks, studied in the context of the European project VirtuWind [2, 7].

The remainder of this paper is organized as follows. In Section II, the background, motivation and related work on Service Function Chaining is presented, highlighting security-related aspects. In Section III, a study on the various use cases and associated variations of the proposed scheme is presented. In Section IV, the Reactive Security Framework and its key implementation elements (e.g. Security Services, Controller modules) are presented, while Section V details the performance evaluation results, along with their assessment in light of the actual network performance recorded in an actual, operational wind park. Section VI concludes this work with some discussion and pointers to future work.

2. SERVICE FUNCTION CHAINING

2.1. Background and Motivation

In typical network deployments, the end-to-end traffic of various applications typically must go through several network services (e.g. firewalls, load-balancers, WAN accelerators). It can also be referred to as Service Functions (or L4-L7 Services, or Network Functions, depending on the source/organisation) that are placed along its path. This traditional networking concept and the associated service deployments have a number of constraints and inefficiencies [8], such as:

Topology constraints: network services are highly dependent on a specific network topology, which is hard to update.

Complex configuration and scaling-out: a consequence of topological dependencies, especially when trying to ensure consistent ordering of service functions and/or when symmetric traffic flows are needed this complexity also hinders scaling out the infrastructure.

Constrained high availability: as alternative and/or redundant service functions must typically be placed on the same network location as the primary one.

Inconsistent or inelastic service chains: network administrators have no consistent way to impose and verify the ordering of individual service functions, other than using strict topologies - on the other hand, these topology constraints necessitate that traffic goes through a rigid set of services functions, often imposing unnecessary capacity and latency costs, while changes to this service chain can introduce a significant administrative burden.

Coarse policy enforcement: classification capabilities and the associated policy enforcements mechanisms are of coarse nature, e.g. using topology information.

Coarse traffic selection criteria: as all traffic in a particular network segment typically has to traverse all the service functions along its path.

The above are exacerbated nowadays, with the ubiquitous use of virtual platforms, which necessitates the use of dynamic and flexible service environments. This is even more pronounced in service provider and/or cloud environments, with infrastructures spanning different domains and serving numerous tenants, each with their own requirements. Said tenants may share a subset of

the providers' service functions, and may require dynamic changes to traffic and service function routing, to follow updates to their policies (e.g. security) or Service Level Agreements.

SFC aims to address these issues via a service-specific overlay that creates a service-oriented topology, on top of the existing network topology, thus providing service function interoperability [5]. An SDN-based SFC Architecture, such as the one defined by the Open Networking Foundation [9], can extend this concept, exploiting the flexibility and advanced capabilities of software defined networks, to provide novel and comprehensive solutions for the above-stated presented weaknesses of the legacy networks.

2.2. Terms and definitions

In this subsection, the terms and their definitions, as used in this work, are mentioned. The definitions of SFC terms are described in IETF, Service Function Chaining (SFC) Architecture [5] and SFC environment Security requirements [10]. Key terms include:

Network Service Function: A function that is responsible for specific treatment of received packets.

Service Function Chaining: A service function chain defines an ordered set of abstract service functions and ordering constraints that must be applied to packets and/or frames and/or flows selected as a result of classification.

Service Function Forwarder: A service function forwarder (SFF) is responsible for forwarding traffic to one or more connected service functions according to information carried in the SFC encapsulation, as well as handling traffic coming back from the service function (legacy or virtual).

Service Function Path: The service function path is a constrained specification of where packets assigned to a certain route must go. Any overlay or underlay technology can be used to create service paths (VLAN, ECMP, GRE, VXLAN, etc.).

Service Function Classifier: An entity that classifies traffic flows for service chaining according to classification rules defined in an SFC Policy Table and to mark packets with the corresponding SF Chain Identifier. It can be on a data path, or run as an application on top of a network controller.

SFC Header: A header that is embedded into the flow packet by the SFC Classifier to facilitate the forwarding of flow packets along the service function chain path. This header also allows the transport of metadata to support various service chain related functionality.

2.3. Related Work

Several SFC-related research efforts can be identified in the literature. Nevertheless, a recent survey on the use of SFC [11] reveals a lack of work focusing on security-related applications, and this is a gap that the framework presented herein can cover. In terms of the key technological building blocks, Network Service Headers (NSH) [12] is an approach that involves the introduction of SFC-specific 4-byte headers that include all the information needed (including associated metadata) to reach a policy decision with regard to what service chain the traffic should follow. As part of the relevant IETF efforts, the NSH approach has been extended to define

a new service plane protocol (a dedicated service plane) for the creation of dynamic service chains [13]; this NSH-based SFC approach is adopted in the framework presented herein.

StEERING [14] is an OpenFlow-based alternative that allows for per-subscriber and per-traffic type/application traffic routing to the various service functions, via simple policies propagated from a centralized control point, but does not consider the security-based classification that forms the basis of the work presented here. Researchers have also introduced SIMPLE [15], a policy enforcement layer that focuses on middleware-specific traffic steering and considers the inclusion of legacy service instances into the chain. It is based on monitoring and correlating packet headers before and after they traverse a specific service function, though this leads to a rather complex process (collecting packets for correlation, matching packets with high accuracy etc.).

The chaining of Virtual Network Functions (VNFs) is another aspect examined in the literature, which considers the trend of virtualizing networks and network functions in modern networks. More specifically, NETSI proposes a security management and monitoring specification in NFV that enable active and passive monitoring of the VNF and the SFC as provisioned in the NFV environment [16]. From this perspective, Megraghdam et al. [17] present a formal model for specifying VNF chains and propose a context-free language for denoting VNF compositions. Chain definitions in the work presented here are based on the structured format required by the testbed controller (i.e. ODL), but a formal-based definition could be used if the corresponding module is appropriately extended, provided that the added complexity is justified by the application requirements. Blendin et al. [18], exploit Linux namespaces to create isolated service instances per service chain, allowing one-to-one mapping of users to service instances; nevertheless, such an approach is not necessary in industrial environments, where, typically, the number of users is limited, and the management of multiple service instances can incur a significant administrative burden.

2.4. Security Service Chaining

Security services are a prime example of traditional network service functions that can benefit from the adoption of SFC, especially in the context of SDN networks. Indeed, security functions such as Access Control List (ACL), Segment, Edge and Application Firewalls, Intrusion Detection and/or Intrusion Prevention systems (IDS/IPS) and Deep Packet Inspection (DPI) are some of the principal service functions considered by IETF when presenting SFC use cases pertaining to Data Centers [19] and Mobile Networks [20]. Said IETF studies consider several SFC use cases and highlight the numerous drawbacks of using traditional service provision methods when applying, among others, the security functions. The security services themselves are typically been deployed as monolithic platforms (often hardware-based), installed at fixed locations inside and/or at the edge of trust domains, and being rigid and static, often lacking automatic reconfiguration and customization capabilities. This approach, combined with the typical networks' architectural restrictions mentioned above, increase operational complexity, prohibit dynamic

updates and impose significant (and often unnecessary) performance overheads, as each network packet must be processed by a series of predefined service functions, even when these are redundant [21].

A typical example of an important, and also ubiquitous, security-related function is DPI, whereby packet payloads are matched against a set of predefined patterns. DPI imposes a significant performance overhead, because of the pattern matching mechanisms that are at the core of its operation, and thus largely unavoidable (motivating a wealth of research efforts focusing on improving their performance [22, 23]). Nevertheless, DPI, in one form or another, is part of many network (hardware or software) appliances and middleboxes; some examples can be seen in [24]. As Bremner-Barr et al. [24] have demonstrated, extracting the DPI functionality and providing it as a common service function to various applications (combining and matching DPI patterns from different sources) can result in significant performance gains; their benchmarks, involving a single Snort-based IDS service function, run in Mininet over OpenFlow to emulate an SDN deployment, compared to two separate traditional instances of Snort, showed that the former (i.e. the single DPI service function) performed 67%-86% faster than the latter.

Leveraging the benefits of SDN-based SFC deployments involves reversing this trend for monolithic, "all-in-one", security services, which are now commonplace. This is an approach, brought forward in part because of the advancements in hardware performance, which meant that a single, relatively affordable, hardware platform had enough resources to accomplish multiple tasks simultaneously. Instead, in the context of SFC, the focus is on breaking-up these complex services into dedicated service functions, each providing a single task. This shift is not dissimilar to the emergence of the Microservices [25] as described in [26], software architectural style (i.e. the Microservices Software Architecture, MSA), which moves developers away from the once-dominant paradigm of building entire applications as a monolith (again, leveraging the benefits of more capable hardware - and mature, sophisticated programming tools), towards applications made up from a number smaller services (elastic, resilient, composable, minimal and complete [27]), each of them performing a single function (adopting the "Do one thing and do it well" philosophy).

Some key security mechanisms to be leveraged in a secure industrial infrastructure, and deployed as virtualized network service functions, appear below:

IDS/IPS is a service able to monitor traffic or system activities for suspicious activities or attack violations, also able to prevent malicious attacks if needed (in the case of IPS).

Honeynet is formed by a set of functions (Honeypots), emulating a production network deployment, able to attract and detect attacks, acting as a decoy or dummy target.

Firewall is a service or appliance running within a virtualised environment providing packet filtering. Legacy firewalls (e.g. actual hardware appliances) are also supported and can easily be integrated into the architecture.

DPI is a function for advanced packet filtering (data and header) running at the application layer of OSI reference model. In DPI packet payloads are matched against a set of predefined patterns.

Network Virtualisation, via the use of Virtual eXtensible Local Area Network (VXLAN [28]), a VLAN-like encapsulation technique to encapsulate MAC-based OSI layer 2 Ethernet frames within layer 4 UDP packets, brings the scalability and isolation benefits needed in virtualised computing environments.

Access Control Lists are used at the entry of the wind park domain to route traffic to the appropriate isolated virtual networks and the corresponding security service functions.

Packet inspectors to detect malformed packets or malicious activity (IPFiX, DDoS)

Secure communication protocols with packet encapsulation services (e.g. IPSec)

Other than the ones employed above, other Service Functions could be included in a real deployment, such as HTTP header enrichment functions, TCP optimisers, Resource Signalling, etc.

3. USE CASES OF SECURITY SERVICE FUNCTION CHAINING

Depending on the focused aspect which is of relevance for each deployment of the proposed service function chaining -enabled framework, some mainly security-focused use cases flavors (sub -use cases) are identified. Each of them is described in the following subsections in more detail.

3.1. Per Tenant Type Classification

One of the main project objectives in industrial networks is faster service provisioning. The time to provision the service is foreseen to be reduced from several days to several minutes. The concept of Service Function Chaining (SFC) has already shown promising results in enabling the faster time-to-market for the new services in the domain of telecom operators. This also implies the potential to reduce CAPEX and OPEX, especially for short lived service. In the context of next generation industrial networks, one of the promised services is the possibility to instantiate virtual tenant networks (VTN) on demand. The purpose of virtual tenant networks is to setup virtual networks which contains a set of functionalities on the same physical hardware platform and network architecture. One objective is that different VTNs are not influencing each other. From industrial networks perspective the tenants are related to different stakeholders in a wind park, shown in the Figure 1. The stakeholders may include: Wind farm operator (WFO), Transmission system grid operator (TSO), Original Equipment Manufacturer (OEM), IoT device vendors, Wind farm owner and SCADA application.

Each of the stakeholders shown in Figure 1 can have different requirements and constraints for setting up a dedicated VTN for their purposes. The different requirements and constraints can lead to different VTN flavors when talking about the VTN setup and network configuration. Service function chaining could be exploited to instantiate VTN, but beyond that aspect, each tenant should be able to add functions on demand in their network, and control the traffic that is allowed within their own VTN. In the above context, two different flavors




Fig. 1. Tenants in Wind Park

of Tenant-based chain classification can be envisioned; presented below.

3.1.1. Security Services on-demand Tenant Use Case

As an example, let's assume that equipment vendor (e.g. Siemens) requests a VTN to inspect the turbines of wind park operator (e.g. EON). Siemens would request a slice of the operator's network allowing it to access (only) the turbine configuration and log files, with certain quality of service (e.g. high availability). Additionally, Siemens must be able to create the list of the engineers and the technicians who are authorized to do a troubleshooting on turbines of a given wind park. Siemens and/or EON would also like to make sure that members of VTN are performing only the authorized tasks (upgrade, traces, etc.). Hence, tenant-specific security components should be added to a security function chain of the wind park. Example of tenant-specific functions are: ACL and DPI. The other security functions, like firewall and IDS might be still shared with other VTNs in the wind park.

In another related scenario, the flexibility of function chaining could allow tenants to change the deployed security mechanisms dynamically. Thus, for example, using tenant-based classification, Tenant 2 could only be using a firewall protection, while Tenant 1 could be using a firewall and an IDS appliance. During operation, the preferences of Tenant 2 could be updated (Tenant 2'), to also necessitate the presence of a Honeypot or Honeynet, triggering the corresponding update to his/her function chain. This is depicted in Figure 2a, which presents an example of such a setup with the following chain definitions:

Chain 1 - Tenant 1: Firewall -> IDS -> Output
 Chain 2 - Tenant 2: Firewall -> Output
 Chain 3 - Tenant 2' (after update): Firewall -> Honeypot/Honeynet -> Output

3.1.2. Industrial Internet of Things Tenant Use Case

Considering the Industrial Internet of Things (IIoT) use case of wind parks deployment, and in the context of

Tenant-based classification, SFC could also be exploited at the application level, in order to provide dynamic, real-time access to the required data of the IIoT sensors. Therefore, depending on tenant's requirements/agreement etc., each of them get presented with a different subset of all parameters monitored by IIoT sensors, even though all tenants will reach the same resource (e.g. web interface on backend monitoring server). This could be achieved by a simple HTTP header enrichment service running on each of the Service Functions, with each of these services adding the corresponding subset of sensed data into the final web pages that the tenants will see on their web browsers. An example of the above concept is depicted in Figure 2b, whereby the following chains are defined:

Chain 1 - IIoT Tenant 1: Sensors' Set 1 -> Output
 Chain 2 - IIoT Tenant 2: Sensors' Set 1 -> Sensors' Set 2 -> Output
 Chain 3 - IIoT Tenant 3: Sensors' Set 1 -> Sensors' Set 2 -> Sensors' Set 3 -> Output

3.2. Per Application Type Classification

This variation of the SFC use cases classifies traffic based on the originating application. Thus, after a stage of Deep Packet Inspection, the Application is identified and the corresponding chain is assigned. An example of chains tailored to specific applications could include forwarding SCADA traffic to a SCADA-specific IDS, and a generic IDS for other traffic, thus limiting the delay imposed on the SCADA traffic by the IDS (as it depends on the number of rules/patterns in the IDS's database, which could be significantly lower in the case of an IDS which only has SCADA-specific rules installed). Another example could be having video surveillance traffic go through Firewall and a Rate Limiter, to lower the transmission rate, respecting QoS requirements. Thus, potential chains in this case (as depicted in Figure 3) could include:

Chain 1 - Unknown application: Firewall -> DPI -> Output
 Chain 2 - Video surveillance: Firewall -> Rate Limiter -> Output
 Chain 3 - Alarm monitoring: Output
 Chain 4 - SCADA: Firewall -> SCADA IDS -> Output
 Chain 5 - Other application: Firewall -> IDS -> Output




Figure 2. SFC - Per Tenant Classification Example




Fig. 3. SFC - Per Application Classification Example

3.3. Per Traffic Type Classification

This use case includes a security SFC-based enhancement, for both intra- and inter- domain deployments, with the ability to forward traffic based on its security classification (e.g. unknown/malicious/legitimate), following predefined Service Function Paths for each traffic type. This type of classification opens up various possibilities for the integration of advanced malicious traffic detection techniques (e.g. exploiting machine learning). As an example, let us assume that a data packet enters the intra-domain wind park deployment. Based on its classification (from the categories listed above), the traffic will be directed to one of three different paths as depicted in Figure 4a. The aim for this process is to route unknown/suspicious traffic via the Intrusion Detection and Deep Packet Inspection Service Functions, in order to classify it (as either legitimate or malicious), thus allowing us to forward it to the windpark or the honeypot, accordingly. Thus, malicious traffic can be isolated in the honeypot, allowing us to track the attacker, identify its purpose and keeping him occupied.

For the inter-domain use case, (Figure 4b), the procedure is similar to the intra-domain scenario. However, a more sophisticated honeypot deployment, such as a Honeynet, can be used as an emulated wind park, having similar services and functions as the original wind park. Moreover, in this case, having acquired the needed tag (as malicious or legitimate) in other parts of the larger wind park deployment, the traffic can avoid going through the same procedures (i.e. Service Functions) again, better highlighting the benefits of SFC in terms of potential performance gains. A core part of this use case is the classifier. The classifier is responsible for classifying and forwarding packets based on predefined rules, exploiting pattern matching and tags found on

the packet headers. The (attached to the SFF) classifier forwards the packets through one of the predefined function chains. In more details, based on the classification of each packet, the traffic can be classified as legitimate, unknown (suspicious) or malicious. Thus, three different chains are defined:

- Chain 1 - Legitimate (known) traffic: Firewall -> Output
- Chain 2 - Suspicious traffic: Firewall -> IDS -> DPI
- Chain 3 - Malicious traffic: Honeybot/Honeynet

4. REACTIVE SECURITY FRAMEWORK IMPLEMENTATION

4.1. Overview

Motivated by the above, this work focuses on providing a security framework to protect critical industrial infrastructures, considering the wind park as a characteristic example, also studying the more complex multi-tenant use case (i.e. a service provider serving multiple tenants; and its evolution, whereby multiple virtual tenant networks have to be established) and the chaining of vital security functions. This work follows closely the standardization efforts of IETF, and the SFC Working Group [29] in specific, building on top of the work of the Open Networking Foundation and the associated OpenDaylight Controller modules, adopting and extending their features. Moreover, special care is given to the security of the SFC mechanisms, e.g. by guaranteeing the integrity of SFC-related data added to the packets for identifying the service functions chains, and by ensuring that no sensitive SFC data (and the associated metadata), crosses different SFC domains, or legacy networks, unprotected.




Fig. 4. SFC - Per Traffic Type Classification Example

One of the goals of this effort is to provide a secure industrial networking infrastructure, via the associated security mechanisms, such as network monitoring and intrusion detection for industrial SDN networks. To achieve this objective, the security framework presented herein includes network monitoring and intrusion detection for identification of attacks and run-time network adaptation for attack response and mitigation mechanisms. By leveraging security network functions such as Firewalls, IDS, DPI, Honeypots and Honeynets, the framework can create a number of service function chains, to forward traffic based on the type of traffic or running application. The aim of this Service Chaining is to overcome constraints and inefficiencies, as mentioned previously. This can be used to fulfil the target of providing security profiles per application classification based on the originating application, or per tenant classification serving multiple virtual tenant networks with the chaining of vital security functions or, alternatively, per traffic classification, for both intra- and inter- domain deployments, following predefined Service Function Paths for each traffic type.

In contrast to the proactive deployment of specific security mechanisms that are setup and deployed before an attack takes place (typically at the network's design phase), the reactive mechanisms employed here are able to react in real-time to changes in the network as well as the traffic traversing said network, e.g. to automatically mitigate attacks, block malicious entities, route them to specific, dummy network components to allow for enhanced monitoring of their actions or even trigger the deployment of new security functions to help alleviate the effects of an ongoing attack. By leveraging the flexibility of SDN-based deployments and the concept of SFC, a service-specific overlay creates a service-oriented topology, on top of the existing network topology, thus providing service function interoperability.

4.2. Deployed Security Service Functions

The reactive security framework includes a number of different service functions as detailed in the subsections below.

4.2.1. IDS and SCADA IDS

The framework's security mechanisms include continuous network monitoring and intrusion detection for identification of attacks and run-time network adaptation for attack response and mitigation mechanisms. More specifically, IDS instances of Snort [30] are deployed, with scripts to ensure that the most up-to-date rules are constantly active. A database for event monitoring is present, while provisions are made to allow for future extensions to transmit relevant information to a security backend (e.g. for more sophisticated pattern matching). Moreover, a SCADA-specific instance of Snort [31] is also deployed, where SCADA traffic will be routed. This limits the delay imposed on the SCADA traffic by the IDS functionality (a delay that significantly depends on the number of rules/patterns in the IDS's database, which will be significantly lower in the case of the IDS which only has SCADA-specific rules installed).

4.2.2. Honeynet

Network-based honeypots have been widely used to detect attacks and malware. A honeypot is a decoy deployment that can fool attackers into thinking they are hitting a real network whereas in the same time it is used to collect information about the attacker and attack method. A Honeynet is a set of functions, emulating a production network deployment, able to attract and detect attacks, acting as a decoy or dummy target. In the protected wind park network, a Honeynet is deployed, consisting of Honeypots emulating SDN and other network elements, as well as Honeypots emulating the operational systems of the wind park, and more specifically elements such as the SCADA systems and the data historian. Simple Honeypots [32], and SCADA-specific Honeypots [33] are deployed

to emulate the exact network and SCADA system setup present in the SDN-enabled wind park. Moreover, passive Honeypots (Early Warning Intrusion Detection Systems, EWIS, in specific [34]) are also be part of the Honeynet, acting as a network telescope on the production part of the industrial network, to monitor all activity in normally unused parts of the network. Such activity is a good indicator of malicious entities operating on the network (such as an attacker probing/foot-printing the network), thus providing early warning of incoming attacks.

4.2.3. Firewall

A software or hardware firewall instance is also deployed on the wind park's network to implement network perimeter security. This is a software firewall (instance of pfsense [35]), but a hardware (legacy) firewall appliance already present in the industrial network could also be used, or even a virtualized commercial firewall appliance (such as the VM-Series from Palo Alto [36]). The type of firewall, as well as its placement, is irrelevant in the context of the reactive security framework employed to protect the industrial network, as the service plane view of the framework focuses on the type of service and not the underlying technology that is used to offer this service, allowing for the use of any type of firewall, and for its placement in any place on an SDN network deployment.

4.2.4. DPI

In the proposed framework's proof-of-concept implementation, nDPI [37] is employed to implement the DPI function, monitor incoming traffic, and assign it to the (sub-)set of security service functions intended for the corresponding traffic type. Since the default nDPI did not meet the framework's requirements, some changes were made to support the SFC applications. In order to have an up to date view of the SFC-related information (e.g. information about the various chains or information on which chain IDs correspond to reverse chains, i.e. return traffic) was fetched from the ODL controller via constantly running scripts.

In terms of packet processing, nDPI listens for pcap packets generated by the IP stack of the kernel. In this implementation, since the packets were forwarded from SFF, a listener was implemented on TCP port 6633 to handle incoming packets from that port; those packets were complete ethernet (ETH) packets that also had NSH headers. Firstly, the NSH header was extracted to check if the packet was already processed (in which case, the packet chain ID would be that of a reverse chain). If it was already processed, it was simply forwarded to the SFF. If the packet has not been processed before, the developed application encapsulates the ETH packet at an pcap packet and then sends it to the *nDPI engine* for processing. As soon as a response is received from the nDPI engine, the appropriate chain IDs and the appropriate next hops from rendered service chains are fetched from the controller, and then, based on the received information, the NSH headers are generated and the ETH packet is encapsulated appropriately, before being forwarded to the SFF. The ETH generation is based on the response of the *nDPI engine* and the appropriate Service Function Chains in order to support dynamic packet flow change. The response of the *nDPI engine* is based on a set

of rules that the engine has compiled to classify traffic types, and can be extended by writing additional rules (for instance TCP and UDP ports 502 associated with Modbus traffic can be defined as being malicious, if such traffic is not expected in the specific part of the network). The response from the *nDPI engine* classification of each packet is either the protocol/application/framework ID that the above mentioned rules define, or *UNKNOWN* if it could not be determined.

4.3. Service Chain Classification

The per-traffic type classification, as described in subsection 3.3, is adopted for the Reactive Security framework, integrating all the Security Service functions detailed in subsection 4.2, via the following implemented Service Function Chains:

Chain 1 - Unknown Traffic: Traffic that is of unknown type (i.e. cannot be classified based on simple ACL rules that the Classifier has), is routed to the *Firewall* and then to the *DPI* (nDPI), where it is analyzed, classified and its headers are updated appropriately, then being assigned to the appropriate Chain (Chains 2 to 4).

Chain 2 - SCADA Traffic: SCADA traffic is routed through the *Firewall* and then through the *SCADA IDS*, which features only SCADA rules, to minimize the performance impact, before being forwarded to its intended destination.

Chain 3 - Legitimate Traffic: Other (non-SCADA) legitimate traffic, is routed through the *Firewall* and then through the generic *IDS*, before being forwarded to the intended destination (in this case, the Data Historian).

Chain 4 - Malicious Traffic: Traffic tagged as malicious (e.g. nmap port scanning), either by the Classifier or by the DPI functionality, is routed through the Firewall and then to the appropriate part of the honeynet; the latter can either be the SCADA Honeypot (if its original target was a SCADA system), the generic Honeypot (if its original target was an SDN device or a production system such as the Data Historian) or the passive EWIS honeypot (if the original target was some unused address, indicating malicious probing/footprinting of the network).

Important parts of the implementation of this functionality are the Classifier and the DPI service function. When there is no previous acquired knowledge about the packet's classification (i.e. no tag on the packet header), the classifier will assign the packet to the Unknown chain (*Chain 1*), aiming to detect any malicious activity, assess its impact, and attach the associated tag, to help form the system's response and enhance the attack mitigation effectiveness. The nDPI disassembles the traffic packets, assesses their content and decides on their traffic type. Then, the packet is repackaged, assigning the appropriate headers to allow for its routing through the corresponding service chain. However, even if this chain will protect SDN network from malicious attacks, the procedure will add delay to the transmission. Thus, in the case of packets already carrying a tag classifying it as legitimate, it will only be forwarded to the firewall (via the associated chain, *Chain 2* or *Chain 3*), providing faster packet transmission. Finally, in case of a malicious type of packets, the classifier will forward the packets to the honeypot (or honeynet, depending on the deployment), via *Chain 4*, to isolate and investigate the attack.




Fig. 5. Graphical User Interface for real-time monitoring of the operation of the Reactive Security Framework on the ODL Controller

4.4. Implemented SDN Controller Modules

To implement the above functionality, other than the security service functions themselves (e.g. IDS, Honeypots) that need to be installed and setup appropriately, certain purpose-built modules as well as enhancements to existing SDN controller modules are needed; in this case for the OpenDaylight (ODL) controller.

4.4.1. SFC Manager

In more detail, the SFC Manager controller module exposes a number of interfaces that various components can use to provide and receive information about service chains that need to be built, which tenants want to use them, which destinations are being accessed, what applications the traffic pertains to and about the service instances of the network functions. Each SFC configuration includes a set of service nodes, a set of service functions, a set of service function forwarders, a set of service chains, a set of service paths and a set of configurations for classifiers (ACL/NSH). The SFC Manager aggregates this information, combines it, and sends service chains in commands to the SDN Controller (OpenDaylight [38] and SFC-ODL [39] are used). The SDN controller, in turn, programs the underlying forwarding elements that do the actual packet forwarding. In essence the SDN Controller is converting commands from the high-level SFC language to the low-level flow filters of expressed in the OpenFlow semantics.

The SDN Controller provides an abstraction view of the network topology. This allows the SFC Manager to focus on the chaining itself and not on the internal topology of the Network Controller. This means SFC Manager manages forwarding rules and flow filters on external ports. This significantly simplifies the configuration. However, the SDN Controller does the necessary transformations to put the paths (sequence of service instances where the packet traverses) and filters (associate user based on his profile to its respective service chain) in the forwarding devices (OF-enabled). In particular the job of SFC Manager is to register external ports of the SDN

transport network (which is being used for SFC) and to declare and associate service instances to those external ports. In the windpark case, such service instances may include vFirewall, IDS, DPI, and Honeypot. These services may be composed of one or multiple instances. These may be the physical appliances or virtual machines running in network function virtualization infrastructure (NFVI). At the Management and Control planes, the SFC Manager and the SFC-enabled SDN controllers are responsible for administrating the services chains, i.e. for translating the operator's/tenant's/application's requirements into service chains. At the Data plane, Classifiers are responsible for assigning traffic to the appropriate service chain (based on various criteria, such as its maliciousness or the tenant that it belongs to, assuming tenant identities have already been validated by authentication/authorization components) and Service Forwarders and Proxies (where needed) are responsible for steering traffic accordingly, in order to realize said Service Chains. The Data plane entities are responsible for steering traffic accordingly, to realize the Service Chains. The Classifier assigns traffic to its intended service chain (based on pre-defined criteria) and the Service Function Forwarders steer traffic to the various Service Function Nodes. If the Service Function Nodes are not OpenFlow-speaking or SFC-aware, or are in different domains, SFC Proxies are needed.

4.4.2. Graphical User Interface

To assess and manage the proof of concept implementation of the Reactive Security Framework, a Graphical User Interface (GUI) was developed, as an additional module on the ODL SDN Controller. The GUI displays instantiated VMs/Service Chains and traffic paths, based on the chains seen in the bottom of Figure 5. Based on this classification, SCADA traffic goes to the SCADA IDS and then to its intended SCADA system at the wind park. HTTP traffic goes to normal IDS and then to its intended system at the wind park. Malicious traffic (e.g. nmap port scan) is detected and goes to Honeypot/Honeynet instead of its intended target wind park system. Finally, unknown traffic is routed to DPI for classification, where a modification

in the header of the packer, can forward the traffic to the respective active chain (legitimate, SCADA or malicious).

To preview the topology of the network, nodeJS library [40] was used to present network topology at real-time. Suitable REST interfaces were implemented to import network components such as switches such as forwarders and classifiers, security functions (FW, DPI, IDS and SCADA-IDS) and end-hosts such as windparks, scada, honeypots, etc. Moreover, the condition of service functions with respect to CPU and memory utilization of the various security service functions (i.e. the VMs running said service functions), is also imported automatically by the use of implemented REST interfaces presented also in real-time on a separate table.

4.5. Architecture Sketch - Module Placement

In today's wind parks security-related functionalities and SCADA applications are running on dedicated locations in the network architecture. Nowadays, these locations have to be decided in the network planning phase and are very difficult to change during the lifetime of a wind park. The approach proposed in this work enables future scenarios to add or delete functionalities or applications in the wind park network during runtime. In order to enhance this flexibility in a network, the principles of NFV MANO can be combined with the framework presented here. The expanded architecture of the framework can be aligned to the approach described in ETSI GS NFV 002 [41], as depicted in Figure 6. This enhances the proposed framework with flexible deployment and instantiation of new VNFs and the automated preparation of service functions chains and will be explored in future work.

5. PERFORMANCE EVALUATION

To evaluate the performance of the reactive framework, an experimental testbed was developed. Furthermore, the per traffic classification was evaluated.

5.1. Testbed Setup

The testbed featuring multiple Virtual Machines (VMs) was developed and deployed on a Proxmox Virtualization Environment [42] an open source server virtualization management software, which run on a server system (featuring 2 x Intel Xeon E5-2630 v2 6-core/12-thread CPUs, at 2.6GHz, with 32GB RAM). The following VMs were required in order to implement the described scenario; three different types of virtual instances were created (in parentheses the resources dedicated to each VM):

1x OpenDaylight Controller instance(Boron release (4 CPU cores, 4GB RAM)), **5x Open vSwitch [43] (v2.59) instances:** (2 x Classifiers (4 CPU cores, 1GB RAM), 3 x Service Function Forwarders (4 CPU cores, 1GB RAM)), **4x Security Service Functions:** (1x Firewall instance (4 CPU cores, 2GB RAM), 1x DPI, i.e. the custom nDPI-based implementation (4 CPU cores, 4GB RAM), 1x Snort-based IDS with all generic rules (4 CPU cores, 1GB RAM), 1x Snort-based SCADA IDS, with SCADA-only rules (4 CPU cores, 1GB RAM), **4x End-hosts:** (1x

Emulated Data Historian (4 CPU cores, 1GB RAM), 1x Emulated SCADA system (4 CPU cores, 1GB RAM), 1x Passive EWIS Honeypot (4 CPU cores, 1GB RAM), 1x SCADA Honeypot (4 CPU cores, 1GB RAM))

5.2. Evaluation Methodology & Results

To give a proof of concept of the implemented framework and testbed, a two steps approach is followed. The first step contains the import of suitable templates (Service Functions, Service Function Forwarders, Service Function Classifiers, Service Function Chains and Access Control Lists) in the controller in JSON formats. These templates includes all the required information of the testbed as presented in the previous subsection. Furthermore, the implemented GUI module depicts realtime network traffic monitoring interfaces and functions and Service Function resource monitoring interfaces and functions custom of the imported data, a screenshot of which can be seen in Figure 5. Changes in path for different traffic types are depicted on the GUI, with different colors to differentiate between the active chains at each instance in time; the various options that can be active (depicting real-time traffic flows and their associated chains) appear in Figure 7. Moreover, in Table I, the results of the conducted experiments are presented. Apart from the delay between end-hosts, following the respective service function chain, the delays between end-hosts i) when there is no function in the middle and ii) when all the security functions are used, are also presented. Although, the results of the experiments are related to the location and the distance of the VMs (in this case all are located on the same server) the correlation between the number of functions and the delays is obvious is useful to evaluate the results of the real traces as are presented in the next subsection.

5.3. Analysis of Results - An Operational Wind Park's Network

In this section, network traces captured from an operational wind park (in Brande, Denmark) are analyzed to highlight the specificities of industrial traffic in the context of this application domain [44], and assess the presented framework's performance in this context.

The subject wind park consists of four wind turbines connected in a redundant star topology. These turbines themselves consist of two switches in series, one at the bottom, the other at the top. Connected to these switches are numerous measurements systems, sensors and actuators which communicate with a SCADA server also connected to the star topology (while these connections are currently over wired networks, they are expected to be replaced with wireless links, tailored to industrial environments [45, 46], in the future). A router then ensures the connection between the central switch and the Internet. The Park Control System mainly consists of two parts: the Wind Farm SCADA System and the Wind Farm Grid Control System. While the SCADA server is responsible for the reporting, supervision, acquisition and storage of data from the turbines, the Control System server is responsible for controlling the power output of the different wind turbines and to adapt it to the requirements received from the grid operator. Traces were gathered from three different locations, as detailed below: traffic to/from




Fig. 6. Expanded, NFV-O Managed and ETSI-aligned, Framework Architecture




Fig. 7. Real-time traffic flows as depicted on the Controller's GUI. Activated Service Chains are color-coded.

the Grid Control System server; traffic to/from the SCADA server; traffic flowing through the intra-domain router. The purpose and requirements of the flows recorded have been analyzed thanks to the input of the network engineers maintaining the subject wind park.

5.3.1. Traffic to/from the Grid Control System Server

The data to and from the Grid Control System server was gathered during approximately 715 seconds. Flows to and from the Grid Control server are mostly instantaneous,

Table I. Experimental Results

Type	DPI	FW	IDS	SCADA IDS	End-Hosts	No of Functions	Delay
No Function					Anywhere	0	0,45 ms
Chain 1 - Legitimate		O	O		Historian	2	66,57 ms
Chain 2 - SCADA		O		O	SCADA	2	45,83 ms
Chain 3 - Malicious		O			Honepot	1	13,53 ms
Chain 4 - Unknown -> Legitimate	O	O	O		Historian	3	169,74 ms
Chain 4 - Unknown -> SCADA	O	O		O	SCADA	3	138,92 ms
Chain 4 - Unknown -> Malicious	O	O			Honepot	2	116,72 ms
All Functions	O	O	O	O	Anywhere	4	191,24 ms

especially for TCP, around 1300 instantaneous flows are observed. Most of these TCP flows correspond to exchanges for the verification of wind characteristics. These flows send a very low amount of data (3 frames of 62 bytes for the data, 3 frames of 60 bytes for the ACKs). Hence, though numerous, they consume a very little amount of data rate. These flows have an end-to-end latency requirement of 500 ms. Of interest are the flows of longer duration. Both for TCP and UDP, these connections send data at a constant rate. However, it is observed that these connections also consume a very low amount of data rate. Indeed, the most consuming connection has an average rate of 529kb/s, which is very low in comparison with the available 1Gb/s links. While the TCP flows correspond to database operations (for logging and scheduling) and have an end-to-end latency requirement of 100 ms, the long duration UDP flows correspond to the regulation communications between the Grid Control server and the turbines. These UDP flows are the most critical and have a strict low end-to-end latency requirement of 10 ms and averages rates of 496kb/s, 80kb/s, 40kb/s and 192b/s depending on the flows.

5.3.2. Traffic to/from the SCADA Server

The data to and from the SCADA server was gathered during approximately 1000 seconds. Whereas the traces contain way much connections than for the Grid Control server (around 20.0000), most of them are best-effort inter-domain access flows. The data rate usage is also very low and most of the connections only send a very small amount of data. More specifically, the most consuming TCP connection consumes 363kb/s while the UDP connections are really negligible with 8kb/s for the most consuming one. The UDP connections mostly consist of Network Time Protocol (NTP) and Dynamic Host Configuration Protocol (DHCP) exchanges, as the SCADA server hosts both a NTP and DHCP server. Though small (average rate in the order of hundreds of b/s), these numerous connections are considered critical and have end-to-end latency requirements of 10 ms and 100 ms respectively. A substantial amount (around 2000) of single-packet Simple Network Management Protocol (SNMP) exchanges with end-to-end latency requirements of 500 ms is also recorded. For their part, the TCP connections mostly consist of online data exchange between the SCADA server and the wind turbines. Though critical, these flows only have end-to-end requirements of 100 ms, 250 ms and 500 ms depending on the specific service. A big amount of instantaneous single-packet UDP exchanges between the SCADA server and the turbines can also be observed.

These have a more stringent end-to-end delay requirement of 10 ms.

5.3.3. Traffic Flowing through the Intra-Domain Router

The data flowing through the intra-domain router was gathered during approximately 1500 seconds. While the data rate consumption observed in the Grid Control and SCADA traces was low, the intra-domain router trace shows even fewer data transmissions. Most of the TCP connections are instantaneous. These correspond to best-effort inter-domain database accesses. The only long duration TCP connection with real-time requirements corresponds to an inter-domain database access with an end-to-end real-time requirement of 100 ms. Some short TCP connections can also be observed, which correspond to regulatory exchanges with real-time requirements of 250 ms or 500 ms. Some observed UDP connections with a (nearly) constant rate correspond to critical NTP connections which have a stringent end-to-end latency requirement of 10 ms. Around 400 seconds and 600 seconds, bursts are observed for several connections. These correspond to SNMP exchanges between the devices in the network and an external network management system. These exchanges are not critical and correspond to scheduled jobs for visual display and reporting.

5.3.4. Discussion

The analysis of the traces has shown that industrial networks have very low data rate requirements. While a lot of short bursty flows exist, only a few of them have QoS requirements. Though these requirements can reach hundreds of milliseconds, a couple of critical flows have stringent low latency requirements of the order of tens of milliseconds. It was also observed that the traces contain a lot of best-effort traffic, also bursty (short duration) for most of them. Even though the average data rate is not a critical factor, this shows that proper QoS management is needed to guarantee that the real-time requirements of critical flows are met even during the short periods when bursts of best-effort database or SNMP exchanges are also transmitted on the network.

Considering the performance of the framework, as evaluated above, and in light of the above findings on the actual wind park, several conclusions can be made. An important observation for the results is that the proposed reactive framework is applicable to operational wind parks, as the additional delays are affordable for most critical services (which all have latency requirements of 100ms, 250ms or 500ms). Furthermore, although the emulated

experiments of service SFC provisioning were conducted between VMs on a same server, thus minimizing network delays, the multiple gigabit interconnections present in a wind park all feature low latencies (through over-provisioning), and are, therefore, expected to introduce minimal delays. Finally, the evaluation of the proposed framework in an actual wind park, validates the feasibility of the approach, providing the necessary sophisticated, dynamic and continuous security monitoring required in industrial networks nowadays.

6. CONCLUSIONS AND FUTURE WORK

This work presented an approach to achieve reactive security for SDN/NFV-enabled industrial networks, based on the use of SFC to dynamically chain various security functions, classify traffic and steer traffic accordingly. The proof-of-concept application of this approach led to the development of a reactive security framework modelled on (and deployable to) an actual, operating wind park, allowing continuous monitoring of the industrial network and detailed analysis of potential attacks, thus isolating attackers and enabling the assessment of their level of sophistication (e.g. from script kiddies to state actors). The deployment of this reactive security framework not only enhances the industrial network's security, but also decreases the performance impact of the security functions. The DPI's performance impact is minimised as the traffic only has to go through one DPI instance, and the same can be said for the IDS/IPS functionality, as e.g. SCADA traffic only has to go through a faster-performing, SCADA-specific IDS instance. The performance evaluation of the framework's implementation validates the feasibility of the approach, also considering the current performance and requirements, as aggregated from an actual, operating wind park.

As future work, improvements will be investigated in both the security service functions as well as the implementation of the DPI functionality (essential for traffic-type classification), to minimize the impact of the framework on the network's performance and enable its use in more time-critical industrial applications. Moreover, the framework will be enhanced via the use of an open source NFV Management and Orchestration (MANO) software stack, which, via the definition of the service templates at the MANO, will be responsible for the boot-up of the necessary VMs using a Virtual Infrastructure Management (VIM) software (e.g. OpenStack). In turn, the MANO will be used to program the ODL Controller accordingly, passing the necessary information to the SFC Manager. This will also enable a more accurate monitoring of the Service Functions' resources (e.g. allowing the instantiation of additional VMs when one of the existing functions is overloaded). Moreover, the automated reactivity of the framework will be enhanced with the integration of SDN security patterns [47] on the ODL controller via the development of an associated model and the introduction of an adaptive access control mechanism that will enable the policy-based management of multiple controllers, across domains [48]. Finally, the performance of the individual components and the

framework as a whole will be evaluated in detail, to assess the impact of the proposed mechanisms in the context of the industrial domain and its associated intricacies. For this purpose, a testbed is already being setup on the operating wind park, in Brande, Denmark, where the trace analysis was conducted; this testbed will form the basis for the real-time evaluation of the framework's performance, as well as its behaviour under different attack scenarios.

ACKNOWLEDGEMENT

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671648 (VirtuWind). The authors would like to thank their colleagues from project VirtuWind for their valuable feedback during the design phase of the framework, as well as the network engineers maintaining the subject wind park for their important input in defining the application requirements and analyzing the network traces.

REFERENCES

1. M. Jose da Silva T. Lins and R. A. Rabelo Oliveira. Software-defined networking for industry 4.0. In *20th Advanced International Conference on Telecommunications*, 2016.
2. N. Petroulakis, T. Mahmoodi, V. Kulkarni, A. Roos, P. Vizaretta, K. Abbasik, X. Vilajosana, S. Spirou, A. Matsiuk, and E. Sakic. Virtuwind: Virtual and programmable industrial network prototype deployed in operational wind park, 2016.
3. Cyber-Attack Against Ukrainian Critical Infrastructure. *Alert (IR-ALERT-H-16-056-01)*, 2015.
4. NERC Standard CIP. 007-6-Cyber Security-Systems Security Management. 2013.
5. Service Function Chaining (SFC) Architecture, RFC 7665. 2015.
6. K. Fysarakis, N. E. Petroulakis, A. Roos, K. Abbasi, P. Vizaretta, G. Petropoulos, E. Sakic, G. Spanoudakis, and I. Askoxylakis. A Reactive Security Framework for Operational Wind Parks Using Service Function Chaining. In *ISCC*, 2017.
7. T. Mahmoodi, V. Kulkarni, W. Kellerer, P. Mangan, F. Zeiger, S. Spirou, I. Askoxylakis, X. Vilajosana, H. Joachim Einsiedler, and J. Quittek. VirtuWind: virtual and programmable industrial network prototype deployed in operational wind park. *Transactions on Emerging Telecommunications Technologies*, 27(9):1281–1288, 2016.
8. P. Quinn and T. Nadeau. Problem Statement for Service Function Chaining, apr 2015.
9. L4-L7 Service Function Chaining Solution Architecture. *Open Networking Foundation*, 2015.
10. Sfc environment security requirements. <https://tools.ietf.org/html/draft-mglt-sfc/securityenvironment-req-01>.
11. D.I Bhamare, R. Jain, M. Samaka, and A. Erbad. A survey on service function chaining. *Journal of*

- Network and Computer Applications*, pages 138 – 155, 2016.
12. P. Quinn and J. Guichard. Service function chaining: Creating a service plane via network service headers. *Computer*, 47(11):38–44, 2014.
 13. P. Quinn and U. Elzur. Network Service Header. *Network Working Group, IETF Draft*, 2016.
 14. Y. Zhang, N. Beheshti, L. Beliveau, G. Lefebvre, R. Manghirmalani, R. Mishra, Ri. Patneyt, M. Shirazipour, R. Subrahmaniam, C. Truchan, and M. Tati-pamula. StEERING: A software-defined networking for inline service chaining. In *ICNP*, 2013.
 15. Z. A. Qazi, C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu. SIMPLE-fying middlebox policy enforcement using SDN. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM - SIGCOMM '13*, page 27, New York, New York, USA, 2013. ACM Press.
 16. GS NFV-SEC 013 ETSI. Network functions virtualisation (nfv) release 3; security; security management and monitoring specification. 2017.
 17. S. Mehraghdam, M. Keller, and H. Karl. Specifying and placing chains of virtual network functions. In *2014 IEEE 3rd International Conference on Cloud Networking, CloudNet 2014*, pages 7–13, 2014.
 18. J. Blendin, J. Ruckert, N. Leymann, G. Schyguda, and D. Hausheer. Position paper: Software-defined network service chaining. In *Proceedings - 2014 3rd European Workshop on Software-Defined Networks, EWSDN 2014*, pages 109–114, 2014.
 19. S. Kumar, M. Tufail, S. Majee, C Captari, and S. Homma. Service Function Chaining Use Cases in Data Centers. 2016.
 20. W. Haeffner, J. Napper, M. Stiernerling, D. Lopez, and J. Uttaro. Service Function Chaining Use Cases in Mobile Networks. 2015.
 21. W. John, K. Pentikousis, G. Agapiou, E. Jacob, M. Kind, A. Manzalini, F. Risso, D. Staessens, R. Steinert, and C. Meirosu. Research Directions in Network Service Chaining. In *IEEE SDN for Future Networks and Services (SDN4FNS)*, pages 1–7. IEEE, nov 2013.
 22. H. Liao, C. Richard Lin, Y. Lin, and K. Tung. Intrusion detection system: A comprehensive review. *Journal of Network and Computer Applications*, 36(1):16–24, 2013.
 23. L. Vokorokos, M. Ennert, M. Cajkovský, and J. Radušovský. A Survey of parallel intrusion detection on graphical processors. *Open Computer Science*, 4(4), jan 2014.
 24. A. Bremler-Barr, Y. Harchol, D. Hay, and Y. Koral. Deep Packet inspection as a service. In *10th ACM International Conference on Emerging Networking Experiments and Technologies, CoNEXT 2014*, pages 271–282, 2014.
 25. Microservices a definition of this new architectural term. <http://martinfowler.com/articles/microservices.html>.
 26. J. Thönes. Microservices. *IEEE Software*, 32, 2015.
 27. Microservices five architectural constraints. <http://www.nirmata.com/2015/02/microservices-five-architectural\discretionary{-}{/}{/}{/}constraints/>.
 28. M. Mahalingam, D. Dutt, K. Duda, P. Agarwal, L. Kreeger, T. Sridhar, M. Bursell, and C. Wright. Virtual extensible local area network (vxlan): A framework for overlaying virtualized layer 2 networks over layer 3 networks, no. rfc 7348. Technical report, 2014.
 29. Service function chaining (sfc) working group. <https://datatracker.ietf.org/wg/sfc/charter/>.
 30. Snort. <http://blog.snort.org/2012/01/snort-292-scada-preprocessors.html>.
 31. Snort 2.9.2: Scada preprocessors. <http://www.snort.org>.
 32. Honeyd. <https://github.com/sk4ld/gridpot>.
 33. Scada honeynet project. <http://scadahoneynet.sourceforge.net>.
 34. P. Chatziadam, I. G. Askoxylakis, and A. Fragkiadakis. A network telescope for early warning intrusion detection. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 8533 LNCS, pages 11–22, 2014.
 35. Open source security. <https://pfsense.org/>.
 36. Vm-series: Next-generation security for private and public clouds. <https://www.paloaltonetworks.com/>.
 37. ndpi: Open and extensible lgplv3 deep packet inspection library. <http://www.ntop.org/products/deep-packet-inspection/ndpi/>.
 38. Opendaylight: Open source sdn platform. <https://www.opendaylight.org>.
 39. Odl wiki: Service function chaining. https://wiki.opendaylight.org/view/Service_Function_Chaining:Main.
 40. Nodejs library. <http://www.nodejs.org>.
 41. ETSI Group Specification NFV 002. Network functions virtualisation (nfv); architectural framework.
 42. Proxmox virtual environment. <http://www.proxmox.com>.
 43. Open virtual switch. <http://www.openvswitch.org>.
 44. Project VirtuWind. Deliverable D3.2: Detailed Intra-Domain SDN & NFV Architecture, 2017.
 45. B. Martinez, X. Vilajosana, Il Kim, J. Zhou, P. Tuset-Peiró, A. Xhafa, D. Poissonnier, and X. Lu. I3Mote: An Open Development Platform for the Intelligent Industrial Internet. *Sensors*, 17(5):986, apr 2017.
 46. D. Dujovne, T. Watteyne, X. Vilajosana, and P. Thubert. 6TiSCH: deterministic IP-enabled industrial internet (of things). *IEEE Communications Magazine*, 52(12):36–41, dec 2014.
 47. N. E. Petroulakis, G. Spanoudakis, and I. G. Askoxylakis. Patterns for the design of secure and dependable software defined networks. *Elsevier Computer Networks*, 109:39–49, 2016.
 48. K. Fysarakis, O. Soultatos, C. Manifavas, I. Papaefstathiou, and I. Askoxylakis. Xsacd - cross-domain resource sharing & access control for smart environments. *Future Generation Computer Systems*, 2016.