

A Novel Cyber-Security Framework Leveraging Programmable Capabilities in Digital Services*

Matteo Repetto¹, Alessandro Carrega², and Armend Duzha³

¹ CNR-IMATI, Genoa, Italy
matteo.repetto@cnr.it

² CNIT, Genoa, Italy
alessandro.carrega@cnit.it

³ Maggioli, Santarcangelo di Romagna, Italy
armend.duzha@maggioli.it

Abstract

The introduction of new computing and networking paradigms, which leverage virtualization and service-oriented architectures, has brought far more agility than ever in the creation and concatenation of digital services. Yet it raises new security and privacy concerns that cannot be properly tackled by existing tools and models.

In this paper, we briefly review the main characteristics of emerging digital services, point out open cyber-security challenges, and discuss the need to include cyber-security programmable capabilities in every digital component. We also describe a novel framework for managing such functions and implement multiple security services for complex business chains.

1 Introduction

Data will be the key driver for the digital economy. Novel digital products and services are expected to create, process, share, and consume data and content in a digital continuum, blurring the frontiers between application domains and breaking the current closed silos of information. One of the most valuable keys to boost innovation and create new revenue streams is unprecedented service agility, which means digital services and business chains are expected to emerge and dissolve much faster than traditional value-creating networks. From a technical perspective, this requires a transition from the design of overarching and standalone services to multi-tenant architectures, leveraging autonomicity and dynamic composition through service-oriented and everything-as-a-service models. The major trend in this respect is the interconnection of processes, products, services, and things from multiple vendors on a growing scale [9], to create, process, share, and distribute data and content, as pictorially depicted for an industrial supply chain in Fig. 1.

The software industry has already been progressively introducing new architectures and patterns that bring more agility in the creation and management of new services and products [8, 1], driving towards fully-automated software and environments that evolve and morph during run-time, without the explicit control of software engineers. The dark side effect of this evolution is the risk for large unpredictability, due to non-deterministic, opaque and partially inscrutable service topology. This raises questions about the overall behavior of the system, the location of personal and sensitive data, the sanity of the software, the availability of the whole service, and,

*This work was supported in part by the European Commission, under Grant Agreement no. 833456.
Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

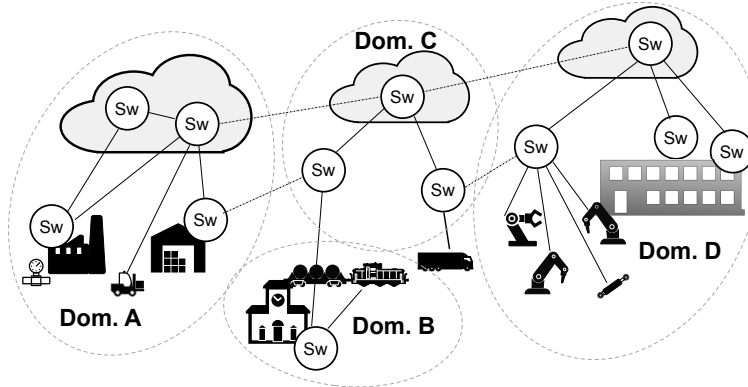


Figure 1: An industrial supply chain creates, processes, shares, and distributes data among multiple actors and ICT infrastructures.

most of all, the ability to perform quick remediation and mitigation actions in case something goes wrong.

New architectures and usage models for building digital services have already revealed the substantial inadequacy of legacy security appliances to effectively protect distributed and heterogeneous systems against cyber-threats [10]. Novel security and privacy models are therefore required, as agile development does not fit the sequential nature of traditional security engineering processes [2].

In this paper, we investigate open cyber-security challenges for digital services that are composed in business value chains. Our main contribution is a conceptual framework where programmable security capabilities are implemented in every digital service and orchestrated by an external logic that takes into account the whole value chain. This approach enables different stakeholders to implement different security services (malware detection, intrusion detection, firewalling and access control, privacy, data tracking, information sharing) tailored to their business, leveraging common components. The framework is currently being designed and developed in the context of the GUARD project, co-funded by the European Commission.

The rest of the paper is organized as follows. Section 2 briefly reviews the main notion and concept of digital services and value chains, pointing out the main challenges from the cyber-security perspective. Sec. 3 highlights the technological gap of existing solutions. Section 4 discusses a novel paradigm, based on the integration of security functions in each digital component that can be orchestrated together to implement coherent security policies over the entire value chain. The preliminary architecture of the orchestration framework and its main logical elements are then described in Sec. 5. Finally, we give our conclusion in Sec. 6.

2 Security challenges for digital value chains

The creation of a digital value chain consists in pipelining several processes, software, and devices, and feeding them with relevant users data and context¹. This has been already possible combining private resources from the same owner; however, the availability of pervasive,

¹See CEF Digital (<https://ec.europa.eu/cefdigital/wiki/display/CEFDIGITAL>) at the policy level, and the FIWARE initiative (<https://www.fiware.org/>) at the technical level.

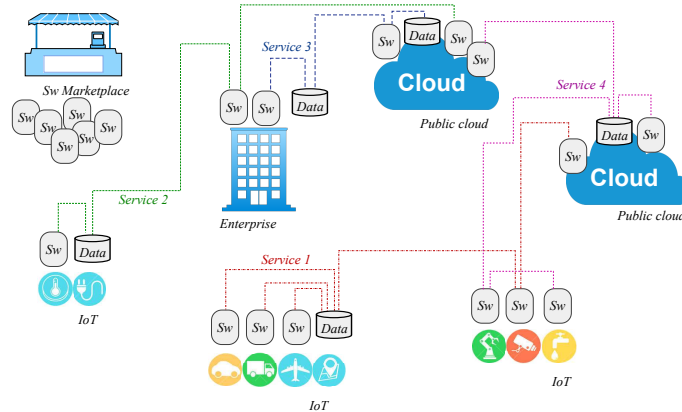


Figure 2: Modern ICT services are generally composed along the three dimensions of infrastructure (legacy enterprise ICT, cloud/fog/edge computing, IoT), software technologies (sw), and data. The definition of a service basically consists in the interconnection of software modules, deployed over heterogeneous infrastructures and domains, and accessing local and remote data.

ubiquitous, and high-performance network connectivity and the uptake of service-oriented architectures have made possible the dynamic composition of heterogeneous digital resources from multiple providers. The set of possible resources entails software, infrastructures, and data: cloud computing (Infrastructure-as-a-Service, Function-as-a-Service, remote storage), Network Functions Virtualization (NFV), Internet of Things (IoT), software repositories, shared authentication services, data sets, environmental and process monitoring are just a few examples of components that are already or will be soon available in the digital market. They are expected to be composed and chained together in elastic topologies, as pictorially shown in Fig. 2.

The composite nature of digital services will lead to complex dynamics, dispersion of data among the multitude of digital objects and infrastructures, unknown or hardly-traceable topologies; in addition, the presence of software components and resource-constrained devices (i.e., smart “things”) in multiple administrative domains makes the application of the security perimeter at each site ineffective, because of the overhead to run complex agents and the difficulty to set-up interoperable and federated mechanisms. Indeed, identity management and access control have already been largely developed and integrated into distributed systems (e.g., cloud computing, IoT, cyber-physical systems), but they can neither guarantee the integrity and dependability of the whole chain over time nor tracking the propagation of private data and sensitive information along industrial, commercial, financial, medical, and other critical processes. Further, a compromised element in a service chain represents a privileged attack vector to affect other systems or the whole chain; it is no coincidence that small devices and smart things are ever-more preferred targets for initiating more complex attacks. Finally, the chain topology and composition are usually unknown to the end user, who cannot easily check whether service owners, security mechanisms (e.g., encryption, integrity), and confidentiality policies are compliant with his own requirements and expectations [13]. The difficulty to assess and certify trustworthiness for multi-domain and multi-tenancy services is currently one main hindrance for commercialization of new products and services with tight constraints on safety and dependability. For instance, in the automotive sector, car manufactures are reluctant to integrate external components (such as road sensors, cloud computing services) into their systems for autonomous driving, just because they have no means to trust information and data.

Just to make a very simple yet illustrative example, let's consider a very conceptual electricity supply chain made of a power plant and a distribution operator. To correctly operate the system, the power injected by the power plant into the grid must balance the current load. We can imagine the indication of the current load as a digital service, which fed the power plant control system. In case the former get compromised, the latter may be driven away from the equilibrium point, with catastrophic consequences for the safety of the electrical grid and the attached users.

The need for novel security paradigms has already been widely recognized, leading to a strong commercial interest in new cyber-security tools that could effectively tackle multi-tenancy, cross-domain interactions, dynamicity, growing complexity of attacks, etc. In particular, the main technological challenges identified so far include: i) situational awareness and its correct and tailored representation to humans; ii) the ability to detect threats and to avoid their propagation in complex and partially unknown business chains; iii) confidence in the trustworthiness and reliability of mutual relationships with other parties involved in the business chains; iv) the ability to share, collect and correlate security-related data from heterogeneous multi-domain and multi-tenancy systems, without disclosing any confidential information; v) the application of analytics by machine learning and other artificial intelligence paradigms in distributed systems.

3 Limitations of existing security models

The cyber-security industry is already tackling the new technological and threat landscape. Many vendors are now delivering integrated solutions for the enterprise, pure and hybrid cloud, industrial devices and networks, and security analytics. Despite the promising trend, we argue that the substantial lack of openness and clear specifications is a major hindrance for achieving interoperability between products from different vendors, which is necessary to align security processes with latest service orchestration and management paradigms (see, for example, the Interface to Network Security Functions (I2NSF) architecture from IETF [3]).

The path towards more effective and efficient security models for digital value chains necessarily requires to go beyond some technical limitations still present today:

- rigidity of the architecture, due to network segmentation and the difficulty to replace or update security appliances;
- inefficient operation, due to the need of bouncing traffic across different security appliances, with partially replicated monitoring and inspection tasks;
- flaws in security appliances, which often represent additional vulnerabilities (see Fig. 3);
- limited visibility, often restricted to a single device or subsystem;
- memory-intensive stateful examination methods, which do not scale to the typical workload of modern services;
- low degree of automation, leaving humans the burden of many repetitive processes.

Today's approaches to check systems during design-time are inapplicable to dynamic service chains given that: a) each service behaves as a black-box, thus it is not necessarily open to inspection; b) service components may be replaced at runtime; and c) the interaction among different services leads to non-deterministic outcomes. Relying on security mechanisms at the infrastructure layer, often implemented in the hypervisor, is no more convenient because of the




Figure 3: Reported vulnerabilities for main vendors of security appliances since 2016. Source: NIST National Vulnerabilities Database.

limited inspection that is allowed by privacy rules, and the difficulty in implementing federation and collaboration mechanisms between providers at both the technical and procedural level [6, 4].

4 Leveraging security capabilities in digital services

A new breed of cyber-security paradigms is therefore needed to protect digital services along the different dimensions. The main challenges in this context can be briefly summarized as follows:

- i) to increase the information base for analysis and detection, while preserving privacy;
- ii) to improve the detection capability by data correlation between domains and sources;
- iii) to verify reliability and dependability by formal methods that take into account configuration and trust properties of the whole chain;
- iv) to increase awareness by better propagation of knowledge to the humans in the loop.

Current challenges and emerging trends are all suggesting the need to evolve from discrete cyber-security appliances to integrated, pervasive and capillary systems, which decouple distributed context monitoring from (logically) centralized detection logic. Key enabler for this evolution will be the architectural separation between analysis and data sources, mediated by proper abstraction; this paradigm will result in an open, modular, pluggable, extendable, and scalable security framework.

From a purely architectural perspective, most cyber-security appliances have been traditionally designed to protect the physical infrastructure, not the services implemented on top of it. This is manifest when considering recent solutions for the cloud, where distributed firewalls, antivirus, intrusion detection systems, and identity/privacy management are often implemented in the hypervisor layer to provide security services to all tenants. Chasing more efficiency, the next evolutionary step is a service-centric architecture that removes the need for legacy security appliances, embeds programmable security capabilities into each software element, and orchestrates them by a common security manager that (logically) centralizes the detection processes. Just like an operating system exposes computational, data and communication resources to applications via Application Programming Interfaces (APIs), so future digital services should




Figure 4: A forward-looking perspective for the evolution of cyber-security architectures for digital services.

include APIs that expose security properties, including configurations, monitoring, inspection, and enforcement rules. The vision is a multi-layer framework (see Fig. 4) that distills trust, privacy, and situational awareness from a large set of data (application logs, security events, packet inspection, network statistics, presence of private and sensitive data, security configurations, etc.), providing tailored messages to different actors (technical, legal and management staff, final users).

5 A framework to implement security services

The framework being developed by the GUARD project includes four main macro-blocks (see Fig. 5):

- *local agents*, concerning the definition of embedded security capabilities that are available within each service, encompassing both inspection and enforcement tasks;
- *security manager*, which is the smart component where a set of security services work in a semi-autonomous manner under the control of user-defined policies;
- *identity management*, to manage certifications, authentication, authorization, and access control in a multi-tenancy environment;
- *user interface*, which allows the definition of custom policies, the delivery of relevant security notifications, and the interaction with other domains, e.g., for sharing Cyber-Threat Intelligence (CTI).

5.1 Local agents

Local agents provides basic security capabilities in each digital resource, including:

- *inspection*: collection of data, events, measurements from heterogeneous sources (application logs, system calls, network traffic), collectively indicated as the “security context,” which can be used to detect attacks and identify new threats;

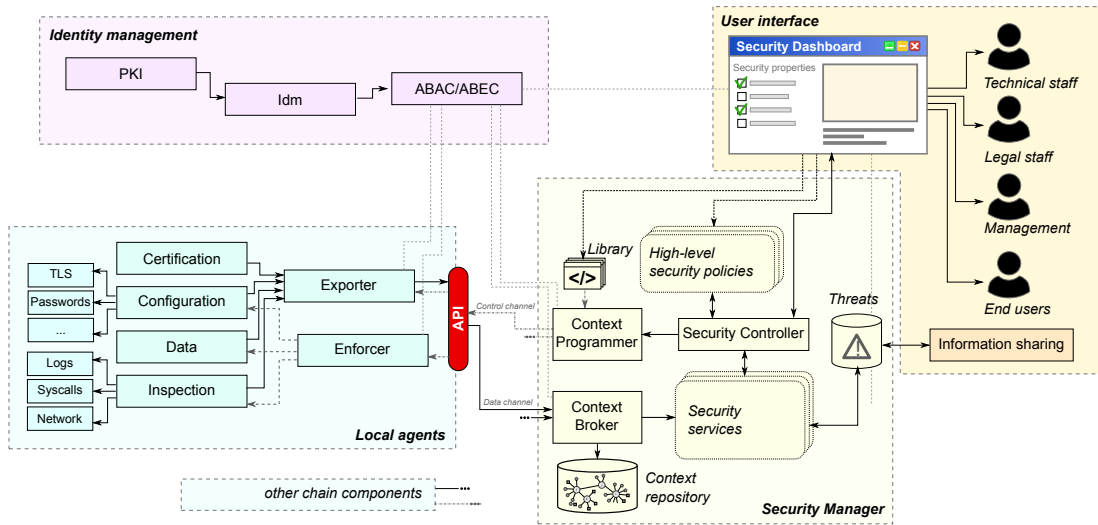


Figure 5: Reference architecture for next-generation cyber-security frameworks for digital value chains.

- *tracking data* belonging to users through metadata, with explicit identification of personal and sensitive information that may raise privacy issues;
- *configuration analysis*, so to report incorrect, faulty, or weak settings: lack of encryption, weak or blank passwords, unnecessary network sockets in listen state, outdated or buggy software versions, etc.;
- *certification* of the origin and integrity of the software component, identity of the vendor/seller, etc.

An exporter function gives access to security-related data and setting. An enforcer function applies security policies: packet classification and filtering, removal of private and/or sensitive data, configuration changes. Enforcement will also cover data protection, by ensuring they are accessed, shared, and exported according to their owner policies in terms of data minimization, purpose limitation, integrity, and confidentiality.

Local agents are expected to be implemented by a heterogeneous set of technologies. In general, they should have small footprint on service execution, so they are required to be efficient and to not expand the attack surface. From a research perspective, the real challenge is *programmability*. Beyond plain *configurability* at run-time (e.g., to adjust the verbosity of logs, frequency of sampling, and other tuneable parameters), programmability also implies the capability to inject inspection, aggregation, pre-processing, and enforcement tasks on-the-fly, without the need for full or partial re-design of the whole system or some of its components. Programming models should target lightweight tasks, to not overwhelm resource-constrained devices, and execution in safe sandboxes, to limit damages from compromised code. Examples of programmable agents include the extended Berkeley Packet Filter (eBPF)² and Logstash³, where custom pipelines can be created to monitor, inspect, and aggregate data.

²<https://www.kernel.org/doc/html/latest/bpf/index.html>.

³<https://www.elastic.co/products/logstash>.

The full set of security capabilities and properties are exposed by a suitable interface. It may be considered as part of the management interface in web services or service-oriented architectures, and will be used to create the security context and enforce security policies.

5.2 Security manager

The Security manager is the most valuable and innovative component in the proposed architecture. It is responsible for orchestrating security capabilities of every digital component and implementing security services, according to the overall objectives and behavior described by high-level user policies. As shown in Fig. 5, multiple logical components are required to implement the Security manager.

Context Broker – The Context Broker collects data from local agents (data channel) and manages the heterogeneity of data sources and protocols, and exposes a common security model to the other components in the run-time subsystem for discovering and accessing the security context available from the execution environment. The definition of a security context model is necessary to know what could be retrieved (i.e., capabilities) and what is currently available, how often, with each granularity (i.e., configuration). The Context Broker hides the heterogeneity and asynchrony of the sources, dispatches the context to multiple consumers, organizes historical data, and provides simple querying and fusion capabilities in data access.

Context Programmer – The Context Programmer is the logical element that offers a homogeneous control interface for configuring and programming different data sources, by implementing the specific protocols (control channel). It changes the reporting behavior by tuning parameters that are characteristics of each local agents (file names, packet filters, frequency and verbosity of reporting, etc.) and offloads lightweight aggregation and processing tasks to each virtual environment. Though the long-term ambition would be the generation of new tasks according to high-level instructions or policies (i.e., a sort of “compiler”), a simplified implementation of the Context Programmer will enable to push pre-defined programs from an internal library. The programs library is a collection of software that can be injected into the programmable hooks present in the execution environment. Different languages can be used by different hooks: Executable and Linkable Format (ELF) binaries, java bytecode, python scripts, P4/eBPF programs. Such programs also include metadata for identification and description, so to be easily referred by the Security Controller. From a security perspective, it is important to formally verify the safety and trustworthiness of the programs. This is implicitly guaranteed, for example, for the eBPF, where the code is executed within an execution sandbox. The Context Programmer has also a context discovery layer: by selectively querying all components involved in the chain, this layer builds the logical topology of the overall service, including the security properties and capabilities of each node. Clearly, this feature relies on a specific capability that must be implemented by a local agent.

Security services – One of the main advantages of collecting heterogeneous security information in a centralized repository is the possibility to implement multiple security appliances in a far more efficient way, i.e. without replicating monitoring and inspection operations. Beyond the mere re-implementation of legacy appliances like Intrusion Detection System (IDS), anti-Denial of Service (DoS), antivirus, and firewall for performance and efficiency matters, the research challenge is a new generation of detection algorithms, arguably by combining rule-based and Machine Learning (ML) methodologies with big data techniques. A preliminary classification of processing and correlation algorithms includes the following issues:

- *attack detection*, which monitors the system behavior to recognize activity patterns that can be associated to known threats and attacks;

- *threat identification*, which identifies anomalies and suspicious activities that deviate from the average system behavior, and tries to define new patterns for unknown and zero-day attacks;
- *data tracking*, which follows the position and transfer of private and sensitive data along the business chain, checks compliance with users privacy policies, and alerts or removes data in case of violations;
- *trust and risk assessment*, which assesses the reliability of the different actors and the resources involved in the business chain, by evaluating the appropriateness of security properties (presence of vendor/software certification, presence of private/sensitive data, configuration settings, availability of security context, use of encryption/integrity mechanisms, availability of specific security features, etc.) to the users policies, and by carrying out evaluation of the risk related to security breaches.

The analysis of the security context from multiple digital resources can greatly enhance the detection capability, especially in case of large multi-vector attacks. The challenge is clearly to merge knowledge without exposing sensitive information to external domains; in this respect, the possibility of local processing provides an effective mechanism. From a practical perspective, however, the real range of algorithms will be limited by the possibility to find an acceptable trade-off between the complexity to implement local inspection and the communication overhead.

Security Controller – The Security Controller automates as much as possible the behavior of the whole framework. It positions between high-level policies and security services, and orchestrates the operation of local agents, according to what already devised in on-going initiatives [5]. Overall, the Security Controller will work according to an Event-Condition-Action (ECA) pattern, where events are triggered by detection or management entities, conditions come from the current context (service topology, security configurations, data and events), and actions entails modifications of the security hooks (monitored data, frequency, granularity, filtering, marking, etc.), re-configuration of the detection algorithms, notifications to users. ECA rules are expressed by policies, which represent the real “smartness” of the Security Controller and encompass both reaction and prevention actions.

High-level security policies – Policies define the behavior of the system. Conceptually, policies do not implement inspection, detection or enforcement tasks, so they do not correspond to any existing security function like IDS, Intrusion Prevention System (IPS), antivirus, Virtual Private Network (VPN). Instead, they represent an additional upper layer for control of security services. Policies are therefore used to automate the response to expected events, avoiding whenever possible repetitive, manual, and error-prone operations done by humans. The simplest way to define behavioral policies is the ECA pattern, which covers a broad range of interesting cases. From a research perspective, the long-term ambition is however the definition of high-level policies in terms of objectives and intents, that could be defined even by non-technical users [7, 5]. The adoption of advanced reasoning models, even based on ML and other forms of artificial intelligence, is clearly a very promising yet challenging target to automate the system behavior. This would open the opportunity for dynamically adapting the response to new threat vectors. In this respect, the historical analysis and correlation of the events and conditions with the effects of the corresponding actions from existing policies or humans would provide useful hints to assess the effectiveness of the latter, so to identify and improve the best control strategies.

5.3 Identity management and access control

The availability of programmable security capabilities in digital resources does not mean they will be publicly available to every external entity. The Identity Management (IdM) module manages the identity of all involved entities as well as the definition of fine-grained access rules. Both the Context Broker and local agents are then expected to enforce such rules. The architecture of this module is largely aligned with best practice for distributed systems, relying on a Public Key Infrastructure (PKI). An IdM component contains a databases that maps the identity of both Users, Local Agents, and any other components belonging to the Security Manager to a specific list of attributes and the Attribute-Based Access Control (ABAC) logic to protect resources.

5.4 User interface

The Security Dashboard is the main management tool used to build situational awareness, to perform reaction and investigation, to share CTI. It can be used to select the set of security services, to visualize anomalies and security events and to pinpoint them in the service topology, to set run time security policies, and to perform manual reaction. However, since bare technical information will be totally useless for most users, the challenge is to deliver tailored informative contents to different levels of the company's structure, to bring awareness to humans and ensure the better understanding. For example, loss or uncertainty in the position of private data should trigger a warning about potential violation of a specific regulation to the legal staff. Any loss of integrity, data, or availability should be reported to the management in terms of potential impact on the overall company business (block of the production, loss of customers, bad reputation). Finally, the main research challenge is the automatic generation of incident reports in standard formats like the Structured Threat Information eXpression (STIX)⁴, their collection in common repositories, and the generation of CTI with attack patterns and threat description [11, 12].

6 Conclusion

The number of recent attacks to IoT devices and cloud services has demonstrated the weak security posture of many digital resources and their potential usage as main vector to more valuable infrastructures and services. The massive digitalization of critical services and infrastructures demands for novel and trustworthy mechanisms for detecting attacks and identifying threats in any link of the value chain.

The GUARD project is tackling this challenge by a ground-breaking framework that leverages the concept of service-oriented architectures and demands for security capabilities to be implemented within each digital block. The GUARD framework is a flexible environment to run multiple security services that orchestrate capabilities of several digital entities involved in a business relationship or value chain. The Consortium is already developing the set of technologies briefly described in this paper, targeting two challenging Use Cases for smart mobility and eHealth; the preliminary software release is planned for next year.

⁴<https://oasis-open.github.io/cti-documentation/>.

7 Acknowledgments

This work was supported in part by the European Commission, under Grant Agreement no. 833456. The design and implementation of the framework described in this paper is a joint effort of the GUARD Consortium. The full list of partners and contributors is available on the project website: <https://guard-project.eu/>.

References

- [1] 5G-PPP. 5g innovations for new business opportunities. Whitepaper, March 2017. [online] Available at <https://5g-ppp.eu/wp-content/uploads/2017/03/5GPPP-brochure-final-web-MWC.pdf>.
- [2] ECSO. European cybersecurity strategic research and innovation agenda (SRIA) for a contractual public-private partnership (cPPP), June 2017. [online] Available at: <https://www.ecs-org.eu/documents/publications/59e615c9dd8f1.pdf>.
- [3] S. Hares, D. Lopez, M. Zarny, C. Jacquenet, R. Kumar, and J. Jeong. Interface to network security functions (I2NSF): Problem statement and use cases. IETF RFC 8192, July 2017.
- [4] Hung-Jen Liao, Chun-Hung Richard Lin, Ying-Chih Lin, and Kuang-Yuan Tung. Intrusion detection system: A comprehensive review. *Journal of Network and Computer Applications*, 36(1):16–24, January 2013.
- [5] D. Lopez, E. Lopez, L. Dunbar, J. Strassner, and R. Kumar. Framework for interface to network security functions. IETF RFC 8329, February 2018.
- [6] Chirag Modi, Dhiren Patel, Bhavesh Borisaniya, Hiren Patel, Avi Patel, and Muttukrishnan Rajarajan. A survey of intrusion detection techniques in cloud. *Journal of Network and Computer Applications*, 36(1):42–57, January 2013.
- [7] D. Montero, M. Yannuzzi, A. L. Shaw, L. Jacquin, A. Pastor, R. Serral-Gracià, A. Lioy, F. Risso, C. Basile, R. Sassu, M. Nemirovsky, F. Ciaccia, M. Georgiades, S. Charalambides, J. Kuusijärvi, and F. Bosco. Virtualized security at the network edge: a user-centric approach. *IEEE Communications Magazine*, 53(4):176–186, April 2015.
- [8] NESSI. Cyber physical systems – opportunities and challenges for software, services, cloud and data. White Paper, October 2015. [online] Available at: http://www.nessi-europe.eu/Files/Private/NESSI_CPS_White_Paper_issue_1.pdf.
- [9] NESSI. Strategic research and innovation agenda, 2017. [online] Available at: http://www.nessi-europe.com/files/NESSI_SRIA_2017_issue_1.pdf.
- [10] R. Rapuzzi and M. Repetto. Building situational awareness for network threats in fog/edge computing: Emerging paradigms beyond the security perimeter model. *Future Generation Computer Systems*, 85:235–249, August 2018.
- [11] Giuseppe Settanni, Florian Skopik, Yegor Shovgenya, Roman Fiedler, Mark Carolan, Damien Conroy, Konstantin Boettinger, Mark Gall, Gerd Brost, Christophe Ponchel, Mirko Haustein, Helmut Kaufmann, Klaus Theuerkauf, and Pia Olli. A collaborative cyber incident management system for European interconnected critical infrastructures. *Elsevier Journal of Information Security and Applications (JISA)*, 34(2):166–182, June 2017.
- [12] Florian Skopik, Giuseppe Settanni, and Roman Fiedler. A problem shared is a problem halved: A survey on the dimensions of collective cyber defense through security information sharing. *Elsevier Computers & Security Journal*, 2016.
- [13] Fulvio Valenza, Cataldo Basile, Daniele Canavese, and Antonio Lioy. Classification and analysis of communication protection policy anomalies. *IEEE/ACM Trans. Netw.*, 25(5):2601–2614, October 2017.