Managing Security for the Cyber-Space - From Smart Monitoring to Automated Configuration

Rémi Badonnel

▶ **To cite this version:**

Rémi Badonnel. Managing Security for the Cyber-Space - From Smart Monitoring to Automated Configuration. Networking and Internet Architecture [cs.NI]. Université de Lorraine (UL), 2022. tel-03606329

HAL Id: tel-03606329
https://inria.hal.science/tel-03606329

Submitted on 30 Dec 2022

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

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

UNIVERSITÉ DE LORRAINE

École doctorale IAEM Lorraine

Managing Security for the Cyber-Space – From Smart Monitoring to Automated Configuration –

MÉMOIRE

présentée et soutenue publiquement le 9 mars 2022

pour l'obtention d'une

Habilitation à Diriger des Recherches de l'Université de Lorraine

(mention informatique)

par

Rémi Badonnel

Composition du jury

Rapporteurs :	Maryline Laurent, Professeur des Universités, Telecom SudParis Ludovic Mé, Adjoint au Directeur Scientifique, Inria Vincent Nicomette, Professeur des Universités, INSA Toulouse
Examinateurs :	Christine Morin, Directrice de Recherche, Inria Filip De Turck, Professeur des Universités, Ghent University Nur Zincir-Heywood, Professeur des Universités, Dalhousie University Stephan Merz, Directeur de Recherche, Inria Pierre-Etienne Moreau, Professeur des Universités, Université de Lorraine
Parrain :	Olivier Festor, Professeur des Universités, Université de Lorraine

Laboratoire Lorrain de Recherche en Informatique et ses Applications — UMR 7503

Mis en page avec la classe thesul.

Acknowledgments

To begin with, I would like to express my thanks to all the distinguished members of this jury, I feel very honoured that you accepted to participate to this habilitation defense.

Many thanks to Maryline Laurent, Ludovic Mé et Vincent Nicomette for their encouraging reports and pertinent comments, as well as to Christine Morin, Filip De Turck, Nur Zincir-Heywood and Stephan Merz for their expertise as examiners.

Special thanks to Pierre-Etienne Moreau for serving as the jury president in addition to being examiner, and for taking care of all the administrative work.

I would like to express my deepest gratitude to Olivier Festor as a scientific mentor for this habilitation, but also for all his great support and confidence since the beginning, both humanly and professionally, as a research team leader and director of TELECOM Nancy.

The contributions presented in this manuscript are the results of several collaborations at national and international levels. I would like to thank all my co-authors, and in particular all the brilliant and dedicated PhD students I have been privileged to work with: Oussema, Martin, Anthea, Maxime, Nicolas, and more recently, Adrien and Mohamed.

I would also like to thank all the permanent and non-permanent members of the MADYNES and RESIST research teams, and in particular Isabelle Chrisment as a research team leader, for all the fruitful research work and rich discussions.

Special thanks also to my PhD supervisors, Andre Schaff and Radu State, for their advices and encouragements since the start of my research career.

Contents

Chapter 1 Introduction **1**

1.1 Research Context . . . 1

1.2 Contributions . . . 2

1.3 Manuscript Organization . . . 5

Chapter 2 Security Monitoring for RPL-based Internet-of-Things **7**

2.1 Introduction . . . 8

2.2 Routing and Monitoring in RPL-based Networks . . . 9

2.3 Taxonomy of Attacks against the RPL Protocol . . . 12

2.4 Impact Assessment of RPL Attacks . . . 15

2.5 Local Mitigation Strategy of DAG Inconsistency Attacks . . . 18

2.6 Security-Oriented Distributed Monitoring Architecture . . . 22

2.7 Conclusions . . . 27

2.8 Related Publications . . . 28

Chapter 3 Vulnerability Management in Autonomic Systems **29**

3.1 Introduction . . . 30

3.2 From Vulnerability Discovery to Remediation . . . 31

3.3 Autonomic Vulnerability Management . . . 34

3.4 Analysis of Past-Hidden Vulnerabilities . . . 39

3.5 Lightweight Assessment based on Probabilistic Scheme . . . 42

3.6 Selection of Corrective Operations using SAT Solving . . . 46

3.7 Conclusions . . . 50

3.8 Related Publications . . . 51

Chapter 4 Software-Defined Security for Distributed Clouds **53**

4.1 Introduction . . . 54

4.2 System Virtualization Models . . . 55

4.3 Security Analysis based on a Reference Architecture . . . 58

4.4 Software-Defined Security Architecture . . . 62
4.5 On-the-Fly Generation of Virtualized Resources . . . 64
4.6 Topology and Orchestration Specification for SDSec . . . 67
4.7 Conclusions . . . 72
4.8 Related Publications . . . 73

Chapter 5 Orchestration of Security Chains in Software-Defined Networks 75
5.1 Introduction . . . 76
5.2 Security Chains in Software-Defined Networks . . . 77
5.3 Overview of the Considered Security Orchestrator . . . 80
5.4 Learning the Networking Behavior of Resources . . . 81
5.5 Synthesis of Security Chains by Inference . . . 86
5.6 Verification Techniques Applied to Security Chains . . . 90
5.7 Conclusions . . . 95
5.8 Related Publications . . . 96

Chapter 6 Conclusions and Perspectives 97
6.1 Conclusions . . . 97
6.2 Research Program . . . 100

Bibliography 105

Glossary 119

List of Figures

1.1 Overview of research activities on security management . . . 2

2.1 Example of a RPL network composed of two instances and three DODAGs. . . . 9
2.2 Taxonomy of attacks against RPL networks. . . . 12
2.3 Total number of control messages experienced by network nodes under different DAG inconsistency attack scenarios. . . . 16
2.4 Total number of loops and inconsistencies for every location of the attacker. . . . 17
2.5 Total control message overhead per node with default, adaptive and dynamic mitigations. $\gamma = 20$ and $\gamma = 25$ are used for the adaptive mitigation. . . . 21
2.6 Example of our monitoring architecture exploiting the RPL multi-instance feature. 22
2.7 Energy consumed by the overhearing mode. . . . 25
2.8 Performance of detection methods with our monitoring architecture. . . . 26

3.1 D^3 classification for vulnerability assessment. . . . 31
3.2 Positioning of vulnerability management with respect to self-management activities. 34
3.3 Vulnerability description with the OVAL language. . . . 35
3.4 Vulnerability awareness high-level architecture. . . . 36
3.5 Distributed vulnerability example. . . . 37
3.6 Assessment of distributed vulnerabilities. . . . 38
3.7 Vulnerability lifecycle events. . . . 39
3.8 High-level imaging and exposure detection process. . . . 41
3.9 Regular vs probabilistic approach. . . . 42
3.10 Impact of the probabilistic strategy. . . . 44
3.11 Ovaldroid implementation prototype. . . . 45
3.12 Change sequence search example. . . . 48
3.13 Performance evaluation. . . . 49

4.1 Virtualization reference architecture. . . . 57
4.2 Synthesis of recommendations with respect to cloud protection. . . . 59
4.3 SDSec architecture overview in a multi-cloud multi-tenant scenario. . . . 62
4.4 Comparison of regular and unikernel VM lifecycles. . . . 65
4.5 Integration of the unikernel generation framework with the SDSec architecture. . 66
4.6 Overview of the TOSCA-oriented SDSec framework for protecting cloud services. 68
4.7 Compared performance of protected unikernel instances. . . . 71

5.1 Considered security chain orchestrator in a SDN infrastructure. . . . 80
5.2 Inferred behavioral automaton for the Pokemon Go Android application. . . . 84
5.3 Simplicity of automata generated with the considered methods. . . . 85

5.4 Accuracy of automata generated for different applications with the considered methods (invarimint in purple color, synoptic in blue color, our approach in green color). The synoptic bar is not plotted for cases where the method was not capable to build the automata. 85
5.5 Synthesis of security chains from behavioral automata. 86
5.6 Illustrative example of simple classification rules. 87
5.7 Extract of elementary security rules and $deploy_r$ predicates. 88
5.8 Illustrative example of a security chain with security functions. 88
5.9 Rewriting rules for translating into Pyretic SDN specification. 89
5.10 Merging and verification of deployable security chains. 90
5.11 Data plane automaton for the toy example. 93

6.1 Research contributions according to the four main axes of our RESIST team. . . 98
6.2 Research program with respect to the four main axes of our RESIST team. . . . 100

1

Introduction

1.1 Research Context

This Habilitation Degree manuscript gives an overview of some of my major research activities performed in the area of network and service management over the past few years.

During my PhD thesis, I started to perform research by investigating and implementing new monitoring methods and techniques for supporting mobile ad-hoc networks [30]. At that time period, these networks were a particularly disruptive networking environment characterized by self-configuring capabilities and spontaneous deployments from mobile devices without requiring any pre-existing fixed infrastructure. The major challenge was to adapt and to make more flexible the network management plane, in order to cope with the properties of these dynamic networks, where nodes may cooperate or not at their will. The research efforts have first focused on building a dedicated information model for ad-hoc networks. We have then reorganized the management plane based on a probabilistic scheme. Instead of considering the whole network, the approach consisted in only selecting nodes that have both a high presence and a strong connectivity with their neighborhood, in order to establish management clusters. Finally, we have adapted management operations, in the context of performance monitoring using filtering techniques, and in the context of fault detection relying on information theory.

At IBM Research, I worked on change management in virtualized infrastructures, in the team of Prof. Joe Hellerstein and Dr. Alexander Keller at Hawthorne Heights, New York. We have proposed the architecture and implementation of a novel workflow-driven provisioning system for application services, such as multi-tiered systems. These services need to be dynamically provisioned to accommodate rapid changes in the workload patterns. This, in turn, requires a highly automated service provisioning process, for which we were able to leverage a general-purpose workflow language, called BPEL4WS, and its execution engine. While the concept of cloud computing was not yet born, we have successfully integrated a workflow-based change management system with a commercial service provisioning system that allowed the execution of automatically generated change plans as well as the monitoring of their execution. In addition to publications, these research efforts have led to an international patent.

My postdoctoral period at the Oslo Metropolitan University took place in the research team of Prof. Mark Burgess, where I worked on new management strategies for autonomic systems. Autonomic computing advocates greater decentralization of autonomy and only weak coupling of components through cooperative communication. It makes traditional server-state and least-connection inapplicable or inefficient. Our efforts on pull-based mechanisms have showed that relaxing the desire for mandatory control of servers using a central controller, and instead allowing

them to cooperate voluntarily through only weak coupling, is not the disadvantage that skeptics imagine; quite the opposite, it has the potential to exceed the performance of a push-based solution, while maintaining better security for each component. In addition to publications, this work has contributed to the integration of new performance metrics into the commonly-used Cfengine configuration management tool, which is developed in this research team, as well as in a dedicated company for the commercial part.

In 2007, I joined as an Associate Professor the TELECOM Nancy School of Engineering in Computer Science, part of the Lorraine INP Collegium, at the University of Lorraine, France, and became a permanent staff member of the RESIST (formerly MADYNES) research team led by Prof. Olivier Festor and then by Prof. Isabelle Chrisment, at the Loria / Inria Nancy Grand Est laboratory. The team activities are focused on network and service management, which is typically organized into five functional domains (FCAPS) which stand for Fault Management, Configuration Management, Accounting Management, Performance Management, and Security Management. Since then, my research efforts, that will be presented in this manuscript, have been centered on novel monitoring and configuration methods and techniques for the functional domain of security management.

1.2 Contributions

Figure 1.1 gives a high-level view on my research activities and their context. These network and service management activities contribute to security management for the cyberspace, in particular the current and new Internet, with the large-scale deployment of the Internet of Things, and the multiplication of services offered by cloud infrastructures. They are structured into three main axes: (1) smart monitoring for low-resource networks, (2) assessment and remediation of vulnerabilities, and (3) automated configuration of virtualized resources.

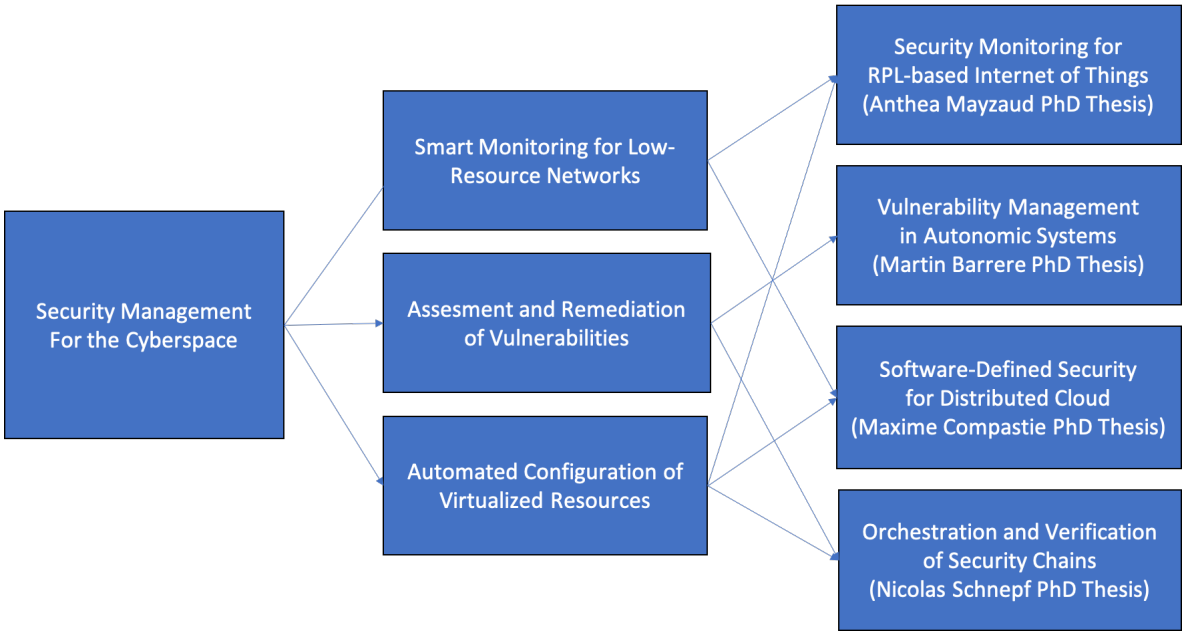

FIGURE 1.1 – Overview of research activities on security management

Smart Monitoring for Low-Resource Networks

A first important axis concerns the investigation of smart monitoring methods capable to cope with low-resource networks, in particular in the context of the Internet of Things. The growing interest for connected objects has resulted in the large scale deployment of Low-power and Lossy Networks (LLN) such as home automation systems. These networks are strongly constrained in terms of resources (memory, power and processing) and communicate using unstable links with high error rates and low throughputs. In this context, existing routing protocols for wired networks and for ad-hoc networks do not cope with all these constraints. More precisely, the IETF RoLL working group has proposed a new routing protocol called RPL based on IPv6 and specially designed for these environments. The RPL protocol is however exposed to a large variety of internal and/or external attacks such as resource consuming attacks, interception or loop building attacks. The deployment of security mechanisms may also be quite expensive in terms of resources. Therefore, LLN networks present new challenges in terms of monitoring and security. During the PhD thesis of Anthéa Mayzaud [130], we proposed a security-oriented monitoring approach for addressing the trade-off between security and cost in the Internet of Things. In a first stage, we assessed security threats faced by these networks. In particular, we identified and classified attacks targeting RPL networks through a dedicated taxonomy. We also quantified the consequences of two major attacks called DAG inconsistency attacks and version number attacks causing over-consumption of node resources. The obtained results showed the importance of addressing them to preserve RPL-based infrastructures. In a second stage, we focused our work on security solutions for RPL-based Internet of Things. We proposed a strategy for addressing DAG inconsistency attacks and evaluated it through experiments. In order to detect more complex attacks such as version number attacks and to complement our node-level approach, we designed a security-oriented distributed monitoring architecture for RPL networks. This solution allowed us to preserve constrained nodes energy by performing monitoring and detection activities on dedicated nodes. We showed the feasibility of our approach by implementing a prototype able to detect both DAG inconsistency and version number attacks. We quantified the performance and the cost of this architecture and the detection modules.

Assessment and Remediation of Vulnerabilities

A second major axis consists in the assessment and remediation of vulnerabilities. The massive deployment of computing devices over disparate interconnected infrastructures has dramatically increased the complexity of network management. Autonomic computing has emerged as a novel paradigm to cope with this challenging reality. By specifying high-level objectives, autonomic computing aims at delegating management activities to the networks themselves. However, when changes are performed by administrators and self-governed entities, vulnerable configurations may be unknowingly introduced. Vulnerabilities constitute the main entry point for security attacks. Hence, self-governed entities unable to protect themselves will eventually get compromised and consequently, they will lose their own autonomic nature. In that context, vulnerability management mechanisms are vital to ensure safe configurations, and with them, the survivability of any autonomic environment. During the PhD thesis of Martin Barrere [34], we targeted the design and development of novel autonomous mechanisms for dealing with vulnerabilities and reducing the exposure to attacks. The contributions concerned autonomic assessment strategies for device-based vulnerabilities and extensions in several dimensions, namely: distributed vulnerabilities (spatial), past hidden vulnerable states (temporal), and mobile security assessment (technological). The spatial dimension permits to cover vulnerabilities that may involve several

devices on a distributed topology. Vulnerability assessment is traditionally performed over individual network devices, independently of each other. Sometimes, however, two or more devices combined together may produce a vulnerable network state that host-based approaches are not able to detect. The temporal dimension focuses on past-hidden vulnerable states. Vulnerability assessment activities usually analyze new security advisories only over current running systems. However, a system compromised in the past by a vulnerability unknown at that moment may still constitute a potential security threat in the present. Indeed, a backdoor installed by an attacker for instance, may remain in the system even though the original vulnerability has been eradicated. The technological dimension aims at reducing the assessment workload in the case of mobile devices, considering a probabilistic cost-efficient technique integrated into a client-server architecture. In addition, we worked on vulnerability remediation methods to autonomously bring networks and systems into secure states. The remediation activity should not generate new vulnerable states on the system. We therefore proposed a formalization of the remediation decision process as a SAT problem, in order to prevent the occurrence of new vulnerabilities when corrective operations are applied.

Automated Configuration of Virtualized Resources

A third axis is dedicated to automating configuration of virtualized resources for supporting security objectives. During the PhD thesis of Maxime Compastié [71], we considered a software-defined security approach for configuring distributed clouds. More specifically, we showed to what extent such programmability facilities can contribute to the protection of distributed cloud services, through the generation of secured unikernel images. These ones are instantiated in the form of lightweight virtual machines, whose attack surface is limited and whose security is driven by a security orchestrator. In that context, we defined a logical architecture supporting the programmability of security mechanisms in a multi-cloud and multi-tenant context. It permits to align and parameterize these mechanisms for cloud services whose resources are spread over several providers and tenants. We then introduced a method for generating secured unikernel images in an on-the-fly manner. It permits to lead to specific and constrained resources, that integrate security mechanisms as soon as the image generation phase. These ones may be built in a reactive or proactive manner, in order to address elasticity requirements. We also extended an orchestration language, so that is possible to generate automatically secured resources, according to different security levels in phase with the orchestration. Complementarily, during the PhD thesis of Nicolas Schnepf [168], we investigated the configuration of security chains. These chains are typically external to the resources to be protected, and may be composed of several security functions, such as firewalls, intrusion detection systems, and data leakage prevention mechanisms. To configure these security chains, it is important to have an adequate model of the patterns that resources (e.g. end user applications) exhibit when accessing the network. We proposed an automated method for learning the networking behavior of resources using algorithms for generating finite state models. These models can be exploited for inferring SDN policies ensuring that applications respect the observed behavior. Such policies can be formally verified and deployed on SDN infrastructures in a dynamic and flexible manner. In particular, our system infers a high-level representation of the security functions, which can be translated into a concrete implementation in the Pyretic language for programming software-defined networks. We showed that the generated chains satisfy several desirable properties such as the absence of black holes and loops, and that they are consistent with the underlying security policy. Further correctness properties of the chains can be verified using our Synaptic checker based on symbolic model checking and SMT solving.

1.3 Manuscript Organization

This manuscript presents a set of contributions in the area of security management addressing the main axes mentioned above. It is structured into four main chapters, complemented by the conclusion chapter.

The first chapter synthesizes research efforts on security monitoring for the Internet-of-Things, performed in the context of the PhD thesis of Anthéa Mayzaud [130], with the design of novel methods and techniques able to cope with the specific properties of these environments, and taking benefits from protocol piggybacking mechanisms.

The second chapter describes research efforts on vulnerability management for autonomic systems, performed in the context of the PhD thesis of Martin Barrère [34]. The presented methods exploit the knowledge provided by configuration vulnerability descriptions, in order to assess the presence of vulnerabilities and select adequate counter-measures.

The third chapter relates to research efforts on software-defined security for distributed clouds, performed in the context of the PhD thesis of Maxime Compastié [71]. The proposed solutions contribute to automate the building and configuration of virtualized resources with a low attack surface in cloud infrastructures.

The fourth chapter is about research efforts on orchestration of security chains, performed in the context of the PhD thesis of Nicolas Schnepf [168]. The approach enables automating the generation and parametrization of these security chains, from an analysis of the networking behaviors of resources.

The last chapter provides conclusions and details research perspectives.

6

Conclusions and Perspectives

6.1 Conclusions

The different contributions detailed in this manuscript are the results of research activities developed in the RESIST research team, which aims at designing, implementing and validating novel models, algorithms and tools to make networked systems elastic and resilient so as to enhance their scalability and security, assuming users, applications and devices whose volume and heterogeneity will continue to increase. The team activities are structured according to four main research axes, namely Monitoring, Experimentation, Analytics and Orchestration, as illustrated on Figure 6.1. Softwarization of networks and data analytics are key enablers to design intelligent methods to orchestrate – i.e. configure in a synchronized and distributed manner – both network and system resources. In particular, intelligent orchestration should leverage relevant data for decision-making using data analytics. Input data reflecting the past, current and even future (predicted) states of the system have to be used for building relevant knowledge. Two approaches can then be pursued to generate knowledge and to validate orchestration decisions [17]. First, a running system can be monitored in vivo. Second, in vitro experimentation in a controlled environment (simulators, emulators and experimental platforms) is helpful to reproduce a running system with a high reliability and under different hypotheses. Monitoring and experimentation are therefore steered and configured through orchestration according to two intertwined loops. As highlighted on the figure, our contributions can be mapped to these team research axes.

Security Monitoring for RPL-based Internet-of-Things

The first contribution is focused on security monitoring for low-power and lossy networks, in the context of the PhD thesis of Anthéa Mayzaud, and mainly relates to the Monitoring and Analytics research axes. Such IoT networks are strongly constrained in terms of resources (memory, power and processing) and communicate using unstable links with high error rates and low throughputs. The IETF RoLL working group has proposed a new routing protocol called RPL based on IPv6 and specially designed for these environments. We have started by assessing security threats targeting the RPL protocol through the identification and classification of attacks and have proposed a dedicated taxonomy. We have analyzed the impact of two RPL specific attacks which are the DAG inconsistency and the version number attacks and showed the importance of addressing them. We have then presented a local strategy to detect and mitigate DAG inconsistency attacks in RPL networks and evaluated its performance and costs. We have designed a security-oriented monitoring architecture in order to complement our node-level approach and address more complex attacks. In a passive and distributed manner, this

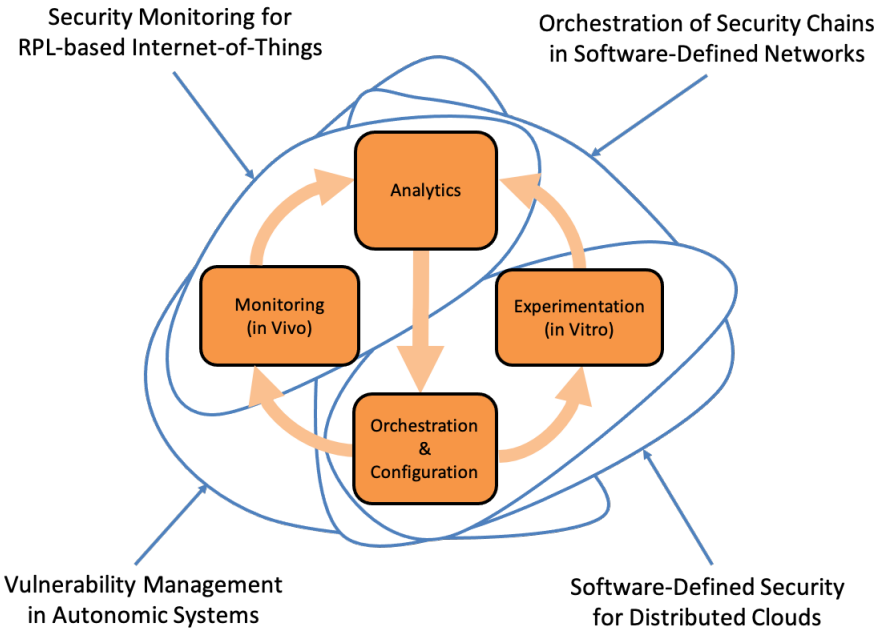

FIGURE 6.1 – Research contributions according to the four main axes of our RESIST team.

solution preserves constrained node resources, by exploiting RPL mechanisms such as the multi-instance feature and by relying on higher-order devices which implement detection modules responsible for identifying the considered security attacks. We have evaluated this architecture through extensive series of experiments and discussed the placement of monitoring nodes in that context. These works performed in the framework of the Flamingo European Network of Excellence, in collaboration with Jacobs University of Bremen, have shown the benefits of protocol piggybacking and dynamic adaptation in order to build a lightweight security-oriented monitoring solution for RPL-based IoT networks.

Vulnerability Management in Autonomic Systems

The second contribution concerns vulnerability management in autonomic systems, in the context of the PhD thesis of Martin Barrère, and mainly relates to the Analytics and Orchestration research axes. Vulnerability management is a major challenge to secure autonomic environments whose changes dynamically operated on their configuration may increase the attack exposure. We have proposed to automate vulnerability assessment by integrating vulnerability descriptions (expressed with the OVAL language) into the autonomic management plane. By translating these security advisories into Cfengine policy rules, autonomic agents deployed across the network become able to analyze their own security exposure and to generate alerts. We have then extended this solution according to three dimensions. We have first addressed distributed vulnerabilities, which correspond to situations where two or more devices under specific conditions may present safe states, but when combined across the network, a vulnerable state arises (spatial dimension). We have then covered the case of past hidden vulnerabilities, by considering an historization of the system configurations from which the vulnerability assessment can be performed to detect past compromissions (temporal dimension). We have also designed a probabilistic solution to lightweight the assessment costs over mobile devices with constrained resources (technological dimension). We have finally worked on vulnerability remediation

mechanisms in order to automate the selection of corrective actions using SAT solving, and to enable collaborative strategies addressing distributed vulnerabilities. These works carried out in the context of the Univerself European project, in collaboration with Alcatel-Lucent Bell Labs, have shown the benefits of automating vulnerability assessment and remediation for maintaining safe configurations in autonomic environments.

Software-Defined Security for Distributed Clouds

The third contribution is about software-defined security for distributed clouds, in the context of the PhD thesis of Maxime Compastié, and mainly relates to the Orchestration and Experimentation research axes. Cloud infrastructures facilitate the provisioning and access to multiple computing resources that require to be efficiently protected, considering the fact that these resources may be distributed over different datacenters, and shared amongst multiple tenants using virtualization technologies. We have first conducted a comparative analysis of virtualization models with regard to cloud protection, and infered several recommendations for our security approach. In particular, we have shown that unikernel-based virtualization provides interesting properties to sanitize the source code and restrict the attack surface, while generating lightweight virtual machines. We have then designed a general software-defined security architecture supporting different abstraction levels to cope with distribution and multi-tenancy, that serves as a basement to our solution, and has been validated based on several use-case scenarios provided by a network operator. We have complemented this architecture with a framework enabling the generation of unikernel cloud resources that comply with security requirements and embed security mechanisms at the earliest stage, as soon as the building of resource images. We have also extended the TOSCA orchestration language to drive such a generation according to different security levels, using our software-defined security architecture. These works developed in the context of the Inria-Orange joint lab, in collaboration with Orange research teams, have shown the benefits of rethinking the security management lifecycle in order to minimize the attack surface, with configuration changes resulting in the systematic rebuilding of cloud resource images in the extreme case.

Orchestration of Security Chains in Software-Defined Networks

The fourth contribution is centered on the orchestration of security chains in software-defined networks, in the context of the PhD thesis of Nicolas Schnepf, and mainly relates to Orchestration and Analytics axes. Network programmability contributes to the flexible building and deployment of security chains for protecting smart devices, such as Android smartphones. It is of major importance to properly configure and verify them to prevent any inconsistencies that could impact on security itself. We have first worked on the automated synthesis of security chains that satisfy by construct correctness properties. For that purpose, we have considered a methodology for profiling the networking behavior of Android applications, and building behavioral models using aggregation and automata learning algorithms, whose performances have been compared in terms of accuracy and simplicity. Based on the obtained models and their properties, we have designed and exploited a rule-based inference system to produce a high-level representation of security chains and their security functions, that are then automatically translated into a concrete implementation deployable in a SDN infrastructure. Complementarily, we have investigated verification and optimization techniques for such security chains that may also be manually specified or updated by network operators. In particular, we have proposed and implemented a security chain checker that supports the rewriting of a set of security chains into different for-

mal specifications that are then interpretable by commonly-used verifiers from SMT solving and model checking areas. These works performed in collaboration with the Veridis INRIA research team have shown the benefits of formal verification to support the building and orchestration of security chains in software-defined networks.

6.2 Research Program

We propose to pursue these research efforts on security automation, according to three main axes, namely (1) ensemble learning methods for smart monitoring, (2) automated security orchestration for composite services, and (3) verified AI-based security management, as illustrated on Figure 6.2. They come within the scope of future network infrastructures, that are characterized by ever-increasing capabilities, in particular in terms of agility, scalability, and automation, as already suggested by the latest deployments of 5G networks and services [19].

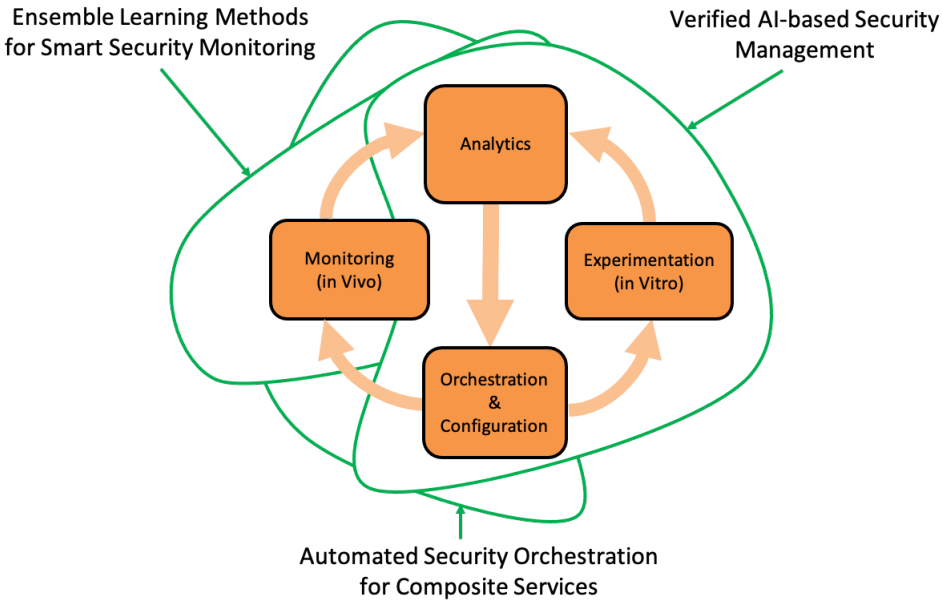

FIGURE 6.2 – Research program with respect to the four main axes of our RESIST team.

Ensemble Learning Methods for Smart Security Monitoring

Security monitoring is challenged by the multiplication and heterogeneity of technologies, protocols and devices that constitute current and future network infrastructures. Learning methods have shown their benefits for building behavioral models from dedicated training datasets, such as network flows and configuration records. These models are then used to identify similarities or deviations characterizing normal behaviors or specific security attacks. For instance, we have already considered them for assessing the network traffic generated by applications running on Android devices, or for parameterizing specific detection methods applied to IoT devices in collaboration with Jacobs University Bremen (Germany) [136]. However, the performance obtained for a given learning method may significantly vary depending on the nature of considered data, which also depends on the strategies that are used for the collection, aggregation and preprocessing of these data, and on the scenarios that are analyzed. In that context, the objective of

our first axis is to investigate ensemble learning methods for enhancing security monitoring. Each learning method has its own advantages and drawbacks, and none of them may outperform the others in all cases. This phenomenon is particularly true with the diversity of monitoring data to be addressed for the applications of future networks, with ever-richer communication technologies. Ensemble learning strategies consist in simultaneously using several learning methods, and combining their results, instead of relying on only one of them. This implies an additional cost to be taken into consideration, but tends to increase the detection performance, and to improve the robustness against adversarial obfuscation techniques.

A typical example for this first axis on ensemble learning methods for smart security monitoring can be given with the case of large-scale heterogenous IoT systems. These latter integrate numerous protocols, platforms and equipments, to support the growing development of smart services, including services for critical and sensitive areas such as industry, transportation and healthcare. The sophistication of security attacks against these systems is increased with the development of advanced persistent threats (APT), such as *Stuxnet* worms against nuclear centrifuges or *Industroyer* malwares against power electric grids, taking the form of multi-step scenarios, that often use to their benefits the complexity of IoT systems, in order to remain as furtive as possible. The detection of these scenarios requires to take into account the attack strategy, the heterogeneity of resources and technologies on which it may rely over time, and the causal relationships amongst its different phases. While different IoT architectures including security modules and features, such as proposed by Carnegie Mellon University (USA) [137] or ARM (UK) [20], have been proposed in the literature, security cannot be guaranteed without failure or only by-design to prevent these attacks, in particular for such evolving ecosystems. Ensemble learning methods have already shown their benefits in many functional areas. For instance, research groups such as Umea University (Sweden) [184], have explored them for supporting fault diagnosis. They also offer promising perspectives to improve attack detection at an early stage in IoT infrastructures, by exploiting complementary learning techniques, such as probabilistic, statistical, proximity-based, and isolation-based methods. However, their usage should not only impact on data analytics, but also on the whole security monitoring process, from the placement and configuration of probes in the network infrastructures to the generation of indicators and alerts, that then serve to orchestrate counter-measures and properly mitigate these advanced attacks.

Automated Security Orchestration for Composite Services

Future infrastructures, leveraged by advances on softwarization, will constitute ever more efficient integration platforms, enabling a higher degree of programmability and automation, in phase with continuous development and delivery strategies, such as those promoted by DevOps and Infrastructure-as-Code initiatives [23]. The growing maturity of orchestration languages already contribute to the building and deployment of composite services. These services typically rely on virtualized resources provided by cloud infrastructures (such as software components and virtual machines) and may be complemented by physical resources (such as connected objects and cyber-physical systems). The orchestration languages permit to specify their structure based on the different resources that compose the service, as well as the relationships that exist amongst them. The objective of our work is to automate the orchestration of security for such composite services, by exploiting and extending the knowledge provided by their specifications. We have already showed the benefits of extending such orchestration languages to drive the building of unikernel-based resources characterized by a low attack surface, in collaboration with Orange [75]. The specified resources and their relationships may provide substantial information

to identify potential vulnerabilities affecting composite services, to enable an adequate placement of security mechanisms, and more generally to improve the resiliency of such services with respect to security attacks. In particular, the relationships include horizontal dependencies, such as two interconnected resources located on different nodes (e.g. a web server and a database server), and vertical dependencies, such as a resource running over another resource (e.g. a web server running over an operating system). This structural information should be taken into account for supporting security automation, and the extension of orchestration languages should be considered for defining different orchestrated security levels, and expressing alternative security mechanisms in order to efficiently adapt to contextual changes. These changes include new security risks that may be identified, collaboratively or not, by cyber threat intelligence, using dedicated tools such as the MISP sharing platform [189] supported by CIRCL (Luxembourg) or the DDoS clearing house [180] developed by SIDN labs (The Netherlands).

A typical example for this second axis on automated security orchestration can be given with the case of cloud composite services and the migration of their resources. The latter may be deployed across different infrastructures owned by one or several cloud provider(s), and are subject to changes over time. This dynamics increases the complexity of management tasks and may lead to potential vulnerabilities that may compromise the resources, or even the whole cloud composite service. In particular, the cold and hot migrations of cloud resources are currently facilitated by recent advances on virtualization techniques, permitting to transfer one or several resource(s) of a cloud composite service from a given provider (or a given infrastructure) to another one. This process is often motivated by performance and cost objectives, with regard to cloud properties, such as scalability, rapid elasticity and on-demand self-service. However, it may impact on the security of cloud composite services and increase their exposure to security attacks. The changes that affect the migrated resources may involuntarily generate vulnerabilities that are exploitable by cybercriminals to cause critical damages, including disclosure of information, data loss and data tampering. It is important to support these migrations with the automated orchestration of adequate security counter-measures. These counter-measures may rely on two categories of security: endogenous mechanisms, such as deploying dedicated security patches, that directly impact internally on the considered resources, or exogenous mechanisms, such as adding new firewall security rules, using different security functions offered externally by cloud providers. Current research efforts typically focus only on one of these categories at a time. For instance, Institut Mines-Telecom (France) [150] have shown how to dynamically generate access control models and policies for different tenant domains, by considering exogenous mechanisms and leveraging network function virtualization (NFV). Some other research groups, such as Ghent University (Belgium) and UFRGS (Brazil) [54], have looked more specifically on the integrity of service function chaining in NFV environments. We believe that orchestration languages provide an interesting and extensible support to partially share security-related information through the usage of trusted third-parties, and to enable an efficient and complementary exploitation of endogenous and exogenous counter-measures, in order to cope with security management issues induced by resource migrations.

Verified AI-based Security Management

Future networks and services will also require further coordination amongst distributed intelligences to enable better operational performance and security amongst networked infrastructures. Artificial intelligence and machine learning are extensively considered for enhancing and automating the different management functional areas, including security management. They help to better identify current and new security threats, and to provide faster responses to security

incidents and attacks, and this being performed over distributed environments. However, their full exploitation is often conditioned by explainability and verification properties, that should ensure the decision-making processes are kept trustful and transparent, and are fully controllable by human administrators and operators. In that context, the objective of the third axis is to bridge the gap between artificial intelligence and verification techniques to support security management automation. We have already investigated management solutions based on verification techniques, such as SAT solving and model checking, to prevent configuration vulnerabilities that may involuntarily occur when changes are operated on the infrastructures in a manual or automated manner [35]. These techniques might also be combined with artificial intelligence methods, in order to guarantee formal verification properties, and improve the decision-making processes. Such an integration may typically be considered for supervised and semi-supervised learning methods, in order to support the preliminary training phase, where the labeling of data might be improved by exploiting the results given by verification techniques. It may also be envisioned for reinforcement learning methods, in order to give verification feedbacks to algorithms that are responsible for exploring different actions and learning from past experiences based on the observations that are performed on the managed system. In the meantime, the considered verification techniques as well as the formal models on which they rely, may in turn benefit from machine learning techniques with respect to their parameterization. Research efforts have already been done at UC Berkeley (USA) [167] to establish the foundations of verified artificial intelligence. They specify the semantic behavior of autonomous industrial systems (e.g. self-driving cars), and then verify that several invariants are guaranteed when AI-based operations are performed. However, they only focus on safety considerations, and do not exploit the knowledge currently provided by security databases and repositories.

A typical example for this third axis on verified AI-based security management can be given with the case of moving target defense (MTD) techniques, that aim at confusing attackers through the reconfiguration of the network infrastructures to be protected. These techniques consist in dynamically changing the available attack surface, by modifying the different resources and parameters of the considered infrastructures, such as the migration of virtual machines, the shuffling of IP addresses, the changes with respect to software product versions, or even the redefinition of functional interfaces. The introduced dynamics impact on the reconnaissance activities, that are performed at the first phase of the cyber kill-chain, by preventing the consolidation of knowledge, regarding the targets and tactics to be considered for performing security attacks, such as accurate identification of entry points and precise software fingerprinting. Methods based on artificial intelligence are currently investigated to support moving target defense strategies, as highlighted by Prof. Gabi Dreo from the CODE Research Institute (Germany) [96]. They automatically determine and schedule the movements to be applied on the infrastructures at different layers. For instance, game theoretic approaches such as developed by Carnegie Mellon University (USA) [129] formalize moving target defense strategies as a two-player game between a defender that continuously shift the system with reconfiguration costs, and an attacker that spends efforts to find new attacks or to try to make past attacks work, the concept of Nash equilibrium being used to establish the defender optimal stationary strategy. The different movements that are applied to the system should not follow any reconfiguration patterns that could be predictable by the attackers. However, this leads to explore new configurations over time that may potentially introduce vulnerabilities on the infrastructure. The changes that are decided by security automation methods based on artificial intelligence, and that may affect different resources distributed over the network, have therefore to be efficiently driven or checked by verification techniques, in order to maintain a minimal attack surface of the considered system.

Bibliography

[1] A Mathematical Programming Language (AMPL). http://www.ampl.com. Last visited on June 2021.

[2] Android Permissions System. https://developer.android.com/guide/topics/security/permissions.html. Last visited on June 2021.

[3] Apache HTTP Server Benchmarking Tool - Apache HTTP Server Version 2.4. http://httpd.apache.org/docs/2.4/en/programs/ab.html. Last visited on June 2021.

[4] Centralized Policy Engine to Enable Multiple OpenStack Deployments for Telco/NFV. https://www.openstack.org/summit/vancouver-2018/summit-schedule/events/21536/. Last visited on June 2021.

[5] Cohttp: Very Lightweight HTTP Server using LWT or Async. https://github.com/mirage/ocaml-cohttp. Last visited on June 2021.

[6] Docker - Build, Ship, and Run Any App, Anywhere. https://www.docker.com/. Last visited on May 2021.

[7] gVisor: Container Runtime Sandbox. https://github.com/google/gvisor. Last visited on June 2021.

[8] IBM Ilog Cplex Optimization Studio. http://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/index.html. Last visited on June 2021.

[9] Linux Containers - LXC - Introduction. https://linuxcontainers.org/fr/lxc/introduction/. Last visited on June 2021.

[10] Mirage Skeleton: Examples of simple MirageOS Applications - Static website TLS. https://github.com/mirage/mirage-skeleton/tree/master/applications/static_website_tls. Last visited on June 2021.

[11] Moon - Security Management Module. https://git.opnfv.org/moon/. Last visited on May 2021.

[12] OCaml Package Manager. http://opam.ocaml.org/. Last visited on June 2021.

[13] Open vSwitch. http://www.openvswitch.org/. Last visited on June 2021.

[14] Oracle VM VirtualBox. https://www.virtualbox.org/. Last visited on June 2021.

[15] QEMU Project. http://www.qemu-project.org/. Last visited on June 2021.

[16] RabbitMQ - Messaging That Just Works. https://www.rabbitmq.com/. Last visited on June 2021.

[17] RESIST Team, Resilience and Elasticity for Security and Scalability of Dynamic Networked Systems, INRIA Activity Report 2020. https://raweb.inria.fr/rapportsactivite/RA2020/resist/. Last visited on October 2021.

[18] XEN Community, Ubuntu Wiki. https://help.ubuntu.com/community/Xen. Last visited on June 2021.

[19] Focus Group on Technologies for Network 2030, International Telecommunication Union (ITU). Network 2030: A Blueprint of Technology, Applications and Market Drivers Towards the Year 2030 and Beyond. https://www.itu.int/en/ITU-T/focusgroups/net2030/Documents/White_Paper.pdf. Last visited on August 2021.

[20] Advanced RISC Machines (ARM). ARM Platform Security Model 1.0. https://developer.arm.com/documentation/den0128/0100/. Last visited on August 2021.

[21] E. S. Al-Shaer and H. H. Hamed. Discovery of Policy Anomalies in Distributed Firewalls. In *Proceedings of the Twenty-third Annual Joint Conference of the IEEE Computer and Communications (IEEE INFOCOM)*, 2004.

[22] S. Arnautov, B. Trach, F. Gregor, T. Knauth, A. Martin, C. Priebe, J. Lind, D. Muthukumaran, D. O'keeffe, and M. Stillwell. SCONE: Secure Linux Containers with Intel SGX. In *Proceedings of the Symposium on Operating Systems Design and Implementation (USENIX OSDI)*, volume 16, pages 689–703.

[23] M. Artac, T. Borovssak, E. Di Nitto, M. Guerriero, and D. A. Tamburri. DevOps: Introducing Infrastructure-as-Code. In *Proc. of the IEEE/ACM 39th International Conference on Software Engineering Companion (IEEE/ACM ICSEC)*, 2017.

[24] L. Atzori, A. Iera, and G. Morabito. The Internet of Things: A Survey. *Elsevier Journal Computer Networks*, 54(15):2787–2805, Oct. 2010.

[25] A. Awad, R. Nebel, R. German, and F. Dressler. On the Need for Passive Monitoring in Sensor Networks. In *In Proceedings of the Conference on Digital System Design Architectures, Methods and Tools (DSD)*, pages 693–699, Sept. 2008.

[26] Z. Ayyub and R. Miao. Simple-fying Middlebox Policy Enforcement using SDN. In *ACM SIGCOMM Computer Communication Review*, 2013.

[27] E. Baccelli, R. Cragie, P. V. der Stok, and A. Brandt. Applicability Statement: The Use of the Routing Protocol for Low-Power and Lossy Networks (RPL) Protocol Suite in Home Automation and Building Control. RFC 7733, Feb. 2016.

[28] M. Backes, S. Bugiel, E. Derr, P. McDaniel, D. Octeau, and S. Weisgerber. On Demystifying the Android Application Framework: Re-visiting Android Permission Specification Analysis. In *Proceedings of the 25th USENIX Security Symposium (USENIX NSDI)*, 2016.

[29] A. Bacs, C. Giuffrida, B. Grill, and H. Bos. Slick: An Intrusion Detection System for Virtualized Storage Devices. In *Proceedings of the 31st Annual ACM Symposium on Applied Computing (ACM SAC)*, pages 2033–2040. ACM Press.

[30] R. Badonnel. *Supervision de Reseaux et Services Ad-Hoc.* Phd Thesis, Henri Poincaré University - Nancy I, France, 2006.

[31] T. Ball, N. Bjørner, A. Gember, S. Itzhaky, A. Karbyshev, M. Sagiv, M. Schapira, and A. Valadarsky. Vericon: Towards Verifying Controller Programs in Software-Defined Networks. In *Proceedings of the 35th ACM SIGPLAN International Conference on Programming Language Design (ACM PLDI)*, pages 282–293, Edinburgh, UK, 2014.

[32] J. Banghart and C. Johnson. The Technical Specification for the Security Content Automation Protocol (SCAP). NIST Special Publication. http://scap.nist.gov/revision/, 2011. Last visited on January 2013.

[33] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield. Xen and the Art of Virtualization. *ACM SIGOPS Operating Systems Review (OSR)*, 37(5):164–177.

[34] M. Barrère. *Vulnerability Management for Safe Configurations in Autonomic Networks and Systems*. Phd Thesis, University of Lorraine, France, 2016.

[35] M. Barrère, R. Badonnel, and O. Festor. A SAT-based Autonomous Strategy for Security Vulnerability Management. In *IEEE/IFIP Network Operations and Management Symposium (IEEE/IFIP NOMS)*, pages 1–9.

[36] M. Barrère, R. Badonnel, and O. Festor. Supporting Vulnerability Awareness in Autonomic Networks and Systems with OVAL. *In Proceedings of the IFIP/IEEE International Conference on Network and Service Management (IFIP/IEEE CNSM)*, Oct. 2011.

[37] M. Barrère, R. Badonnel, and O. Festor. Collaborative Remediation of Configuration Vulnerabilities in Autonomic Networks and Systems. *In Proceedings of the 8th IFIP/IEEE International Conference on Network and Service Management (IFIP/IEEE CNSM)*, Oct. 2012.

[38] M. Barrère, R. Badonnel, and O. Festor. Ovalyzer: an OVAL to Cfengine Translator. In *Proceedings of the IEEE/IFIP Network Operations and Management Symposium (IEEE/IFIP NOMS)*. Ph.D. Student Demo Contest of the IFIP/IEEE Network Operations and Management Symposium (NOMS'12), Apr. 2012.

[39] M. Barrère, R. Badonnel, and O. Festor. Towards the Assessment of Distributed Vulnerabilities in Autonomic Networks and Systems. In *Proceedings of the IEEE/IFIP Network Operations and Management Symposium (IEEE/IFIP NOMS)*, pages 335–342, Apr. 2012.

[40] M. Barrère, R. Badonnel, and O. Festor. Improving Present Security through the Detection of Past Hidden Vulnerable States. In *Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management (IFIP/IEEE IM)*, May 2013.

[41] M. Barrère, R. Badonnel, and O. Festor. Vulnerability Assessment in Autonomic Networks and Services: A Survey. *IEEE Communications Surveys & Tutorials*, PP(99):1–17, 2013.

[42] M. Barrère, G. Betarte, and M. Rodríguez. Towards Machine-assisted Formal Procedures for the Collection of Digital Evidence. In *Proceedings of the 9th Annual International Conference on Privacy, Security and Trust (PST)*, pages 32 –35, July 2011.

[43] M. Barrère, G. Hurel, R. Badonnel, and O. Festor. Increasing Android Security using a Lightweight OVAL-based Vulnerability Assessment Framework. *In Proceedings of the 5th IEEE Symposium on Configuration Analytics and Automation (IEEE SafeConfig)*, Oct. 2012.

[44] M. Barrère, G. Hurel, R. Badonnel, and O. Festor. A Probabilistic Cost-efficient Approach for Mobile Security Assessment. In *Proceedings of the IFIP/IEEE International Conference on Network and Service Management (IFIP/IEEE CNSM)*, Zurich, Switzerland, Oct. 2013.

[45] M. Barrère, G. Hurel, R. Badonnel, and O. Festor. Ovaldroid: an OVAL-based Vulnerability Assessment Framework for Android. In *Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management (IFIP/IEEE IM)*, Ghent, Belgium, May 2013. IEEE.

[46] A. Baumann, M. Peinado, and G. Hunt. Shielding Applications from an Untrusted Cloud with Haven. *ACM Transactions on Computer Systems*, 33(3):8:1–8:26.

[47] R. Beckett. *Network Control Plane Synthesis and Verification*. PhD thesis, University of Princeton, 2018.

[48] C. J. Bernardos, A. Rahman, J.-C. Zuniga, L. M. Contreras, P. A. Aranda, and P. Lynch. Network Virtualization Research Challenges. RFC 8568, Apr. 2019.

[49] I. Beschastnikh, J. Abrahamson, Y. Brun, and M. D. Ernst. Synoptic: Studying Logged Behavior with Inferred Models. In *Proceedings of the ACM SIGSOFT Symposium and the European Conference on Foundations of Software Engineering (ESEC/FSE*, pages 448–451, New York, NY, USA, 2011. ACM.

[50] I. Beschastnikh, Y. Brun, J. Abrahamson, M. D. Ernst, and A. Krishnamurthy. Using Declarative Specification to Improve the Understanding, Extensibility, and Comparison of Model-Inference Algorithms. In *IEEE Transactions on Software Engineering*, volume 41, pages 408–428, 2015.

[51] A. Biere, M. Heule, H. van Maaren, and T. Walsch. *Handbook of Satisfiability.* IOS Press, 2009.

[52] A. Biermann and J. Feldman. On the Synthesis of Finite-State Machines from Samples of Their Behavior. In *IEEE Transactions on Computers*, 1972.

[53] R. Bohme. Vulnerability Markets. What is the Economic Value of a Zero-Day Exploit ? In *Proceedings of the 22nd Chaos Communication Congress*, December 2005.

[54] L. Bondan, T. Wauters, B. Volckaert, F. De Turck, and L. Z. Granville. Anomaly Detection Framework for SFC Integrity in NFV Environments. In *Proceedings of the IEEE International Conference on Network Softwarization (IEEE NetSoft)*, 2017.

[55] C. Bormann, M. Ersue, and A. Keranen. IETF RFC 7228 - Terminology for Constrained-Node Networks. https://datatracker.ietf.org/doc/html/rfc7228, May 2014.

[56] A. Bratterud, A. Walla, H. Haugerud, P. E. Engelstad, and K. Begnum. IncludeOS: A Minimal, Resource Efficient Unikernel for Cloud Services. In *Proceedings of the IEEE International Conference on Cloud Computing Technology and Science (IEEE CloudCom)*, pages 250–257.

[57] M. Burgess and Æ. Frisch. *A System Engineer's Guide to Host Configuration and Maintenance Using Cfengine*, volume 16 of *Short Topics in System Administration.* USENIX Association, 2007.

[58] J. Caballero, Z. Liang, P. Poosankam, and D. Song. Towards Generating High Coverage Vulnerability-Based Signatures with Protocol-Level Constraint-Guided Exploration. In *Proceedings of the 12th International Symposium on Recent Advances in Intrusion Detection (RAID)*, pages 161–181. Springer-Verlag, 2009.

[59] J. Cappos, J. Samuel, S. Baker, and J. H. Hartman. A Look in the Mirror: Attacks on Package Managers. In *Proceedings of the 15th ACM Conference on Computer and Communications Security (ACM CCS)*, pages 565–574. ACM Press.

[60] M. Carpenter, T. Liston, and E. Skoudis. Hiding Virtualization from Attackers and Malware. *IEEE Security Privacy*, 5(3):62–65, 2007.

[61] J. Case, M. Fedor, M. Schoffstall, and J. Davin. Simple Network Management Protocol (SNMP). RFC 1157 (Historic), May 1990.

[62] Cfengine. `http://www.cfengine.com/`. Last visited on November 2021.

[63] Chef. `http://www.getchef.com/chef/`. Last visited on November 2021.

[64] B.-r. Chen, G. Peterson, G. Mainland, and M. Welsh. LiveNet: Using Passive Monitoring to Reconstruct Sensor Network Dynamics. In S. Nikoletseas, B. Chlebus, D. Johnson, and B. Krishnamachari, editors, *Distributed Computing in Sensor Systems*, volume 5067 of *Lecture Notes in Computer Science*, pages 79–98. Springer Berlin Heidelberg, 2008.

[65] D. R. Cheriton and K. J. Duda. A Caching Model of Operating System Kernel Functionality. In *Proceedings of the USENIX Conference on Operating Systems Design and Implementation (USENIX OSDI)*. USENIX Association, 1994.

[66] M. Christodorescu, R. Sailer, D. L. Schales, D. Sgandurra, and D. Zamboni. Cloud Security Is Not (Just) Virtualization Security: a Short Paper. In *Proceedings of the ACM Workshop on Cloud Computing Security (ACM CCSW)*, page 97. ACM Press.

[67] C. J. Chung, P. Khatkar, T. Xing, J. Lee, and D. Huang. NICE: Network intrusion detection and countermeasure selection in virtual network systems. *IEEE Transactions on Dependable and Secure Computing*, 10(4):198–211, 2013.

[68] Cisco Systems. Routing in The Internet of Things – M2M Networks. *BRKSPG-1661*, 2013.

[69] E. M. Clarke, T. A. Henzinger, H. Veith, and R. Bloem, editors. *Handbook of Model Checking*. Springer, 2016.

[70] P. Colp, M. Nanavati, J. Zhu, W. Aiello, G. Coker, T. Deegan, P. Loscocco, and A. Warfield. Breaking Up is Hard to Do: Security and Functionality in a Commodity Hypervisor. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles (ACM SOSP)*, pages 189–202. ACM Press.

[71] M. Compastie. *Software-defined Security for Distributed Clouds*. Phd Thesis, University of Lorraine, France, 2018.

[72] M. Compastiè, R. Badonnel, O. Festor, and R. He. A TOSCA-Oriented Software-Defined Security Approach for Unikernel-Based Protected Clouds. In *Proceedings of the IEEE Conference on Network Softwarization (IEEE NetSoft)*, pages 151–159, 2019.

[73] M. Compastié, R. Badonnel, O. Festor, and R. He. From Virtualization Security Issues to Cloud Protection Opportunities: An In-depth Analysis of System Virtualization Models. *Computers and Security*, 97:101905, 2020.

[74] M. Compastié, R. Badonnel, O. Festor, R. He, and M. Kassi-Lahlou. Towards a Software-Defined Security Framework for Supporting Distributed Cloud. In *Proceedings of the IFIP International Conference on Autonomous Infrastructure, Management and Security (IFIP AIMS)*, pages 47–61. Springer.

[75] M. Compastié, R. Badonnel, O. Festor, R. He, and M. Kassi-Lahlou. Unikernel-based Approach for Software-defined Security in Cloud Infrastructures. In *Proceedings of the IEEE/IFIP Network Operations and Management Symposium (IEEE/IFIP NOMS)*, pages 1–7, Apr. 2018.

[76] V. Corey, C. Peterman, S. Shearin, M. Greenberg, and J. Van Bokkelen. Network Forensics Analysis. *Internet Computing, IEEE*, 6(6):60 – 66, Nov 2002.

[77] J. Criswell, N. Dautenhahn, and V. Adve. Virtual Ghost: Protecting Applications from Hostile Operating Systems. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems (ACM ASPLOS)*, pages 81–96. ACM Press.

[78] CVE, Common Vulnerabilities and Exposures. `http://cve.mitre.org/`. Last visited on November 2021.

[79] CVSS, Common Vulnerability Scoring System. `http://www.first.org/cvss/`. Last visited on November 2021.

[80] N. Dautenhahn, T. Kasampalis, W. Dietz, J. Criswell, and V. Adve. Nested Kernel: An Operating System Architecture for Intra-Kernel Privilege Separation. *SIGARCH Computer Architecture News*, 43(1):191–206.

[81] P. V. der Stok and A. Bierman. CoAP Management Interface. Internet-Draft draft-vanderstok-core-comi-09, Internet Engineering Task Force, mar 2016. Work in Progress.

[82] A Road Map for Digital Forensic Research. In Report From the First Digital Forensic Research Workshop (DFRWS). http://www.dfrws.org/2001/dfrws-rm-final.pdf, August 2001.

[83] J. Dilip and S. Ion. Modeling Middle Boxes. In *IEEE Network: The Magazine of Global Internetworking Archives*, 2008.

[84] S. Dobson, F. Zambonelli, S. Denazis, A. Fernández, D. Gaïti, E. Gelenbe, F. Massacci, P. Nixon, F. Saffre, and N. Schmidt. A Survey of Autonomic Communications. *ACM Transactions on Autonomous and Adaptive Systems*, 1(2):223–259, Dec. 2006.

[85] D. Dong, X. Liao, Y. Liu, C. Shen, and X. Wang. Edge self-monitoring for wireless sensor networks. *IEEE Transactions on Parallel and Distributed Systems*, 22(3):514–527, March 2011.

[86] D. R. Engler, M. F. Kaashoek, and J. O'Toole, Jr. Exokernel: An Operating System Architecture for Application-level Resource Management. *SIGOPS Operating System Reviews*, 29(5):251–266.

[87] R. Enns, M. Bjorklund, A. Bierman, and J. Schönwälder. Network Configuration Protocol (NETCONF). RFC 6241, Oct. 2015.

[88] European Commission. FP7 Flamingo Network of Excellence on Management of the Future Internet. https://cordis.europa.eu/project/id/318488, 2016.

[89] P. Faruki, A. Bharmal, V. Laxmi, V. Ganmoor, M. S. Gaur, M. Conti, and M. Rajarajan. Android Security, a Survey of Issues, Malware Penetrations and Defenses. In *IEEE Communications Surveys & Tutorials*, 2015.

[90] N. Feamster, J. Rexford, and E. Zegura. The Road to SDN, an Intellectual History of Programmable Networks. *SIGCOMM Computer Communication Review*, 44(2):87–98, 2014.

[91] W. Felter, A. Ferreira, R. Rajamony, and J. Rubio. An Updated Performance Comparison of Virtual Machines and Linux Containers. In *Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (IEEE ISPASS)*, pages 171–172.

[92] N. Foster, M. J. Freedman, A. Guha, R. Harrison, N. P. Kata, C. Monsanto, J. Reich, M. Reitblatt, R. Jennifer, C. Schlesinger, A. Story, and D. Walker. Languages for Software-Defined Networks. In *Software Technology Group*, 2016.

[93] N. Foster, M. J. Freedman, R. Harrison, C. Monsanto, and D. Walker. Frenetic, a Network Programming Language. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming (ACM ICFP)*, 2011.

[94] N. Foster, J. McClurg, H. Hojjat, and P. Cerny. Efficient Synthesis of Network Updates. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation (ACM PLDI)*, 2015.

[95] S. Frei, D. Schatzmann, B. Plattner, and B. Trammel. Modelling the Security Ecosystem - The Dynamics of (In)Security. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*, June 2009.

[96] Gabi Dreo. AI-Based Cyber Defence: Two Sides of The Medal, Keynote Talk at the International Conference on Integrated Network Management (IFIP/IEEE IM). 2021.

[97] F. P. Garcia, R. M. C. Andrade, C. T. Oliveira, and J. N. de Souza. EPMOSt: An Energy-Efficient Passive Monitoring System for Wireless Sensor Networks. *Sensors*, 14(6):10804, 2014.

[98] GData. Mobile Malware Report. http://www.gdatasoftware.com, 2019.

[99] D. Gislason. *Zigbee Wireless Networking*. Newnes, USA, PAP edition, 2008.

[100] S. Godik, T. Moses, A. Anderson, B. Parducci, C. Adams, D. Flinn, G. Brose, H. Lockhart, K. Beznosov, M. Kudo, and others. *EXtensible Access Control Markup Language (XACML) version 1.0*. 2003.

[101] A. Goel, K. Po, K. Farhadi, Z. Li, and E. de Lara. The Taser Intrusion Recovery System. *SIGOPS Operating System Reviews*, 39(5):163–176, 2005.

[102] H. Hu, W. Han, G. Joon Ahn, and Z. Zhao. Flowgard : Building Robust Firewalls for Software-Defined Networks. In *Proceedings of the third Workshop on Hot topics in Software-Defined Networking (ACM SIGCOMM)*, 2014.

[103] J. Hui and J. Vasseur. RFC 6553: The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams. https://datatracker.ietf.org/doc/html/rfc6553, Mar. 2012.

[104] G. Hurel, R. Badonnel, A. Lahmadi, and O. Festor. Behavioral and Dynamic Security Functions Chaining for Android Devices. In *Proceedings of the 11th IFIP/IEEE/ACM SIGCOMM International Conference on Network and Service Management (IFIP/IEEE/ACM SIGCOMM CNSM)*, 2015.

[105] G. Hurel, R. Badonnel, A. Lahmadi, and O. Festor. Towards Cloud Based Compositions of Security Functions for Mobile Devices. In *Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management (IFIP/IEEE IM)*, 2015.

[106] IBM. http://www.ibm.com/. Last visited on November 2021.

[107] X. Jiang, X. Wang, and D. Xu. Stealthy Malware Detection Through Vmm-based "Out-of-the-box" Semantic View Reconstruction. In *Proceedings of the 14th ACM Conference on Computer and Communications Security (ACM CCS)*, pages 128–138. ACM Press.

[108] H. Kang, M. Le, and S. Tao. Container and Microservice Driven Design for Cloud Infrastructure DevOps. In *Proceedings of the IEEE International Conference on Cloud Engineering (IEEE IC2E)*, pages 202–211.

[109] M.-Y. Kang, J.-Y. Choi, I. Kang, H. H. Kwak, S. J. Ahn, and M.-K. Shin. *A Verification Method of SDN Firewall Applications*. IEICE Transactions on Communications, 2016.

[110] M. M. H. Khan, L. Luo, C. Huang, and T. Abdelzaher. SNTS: Sensor Network Troubleshooting Suite. In *Proceedings of the 3rd IEEE International Conference on Distributed Computing in Sensor Systems*, DCOSS'07, pages 142–157, Berlin, Heidelberg, 2007. Springer-Verlag.

[111] A. Khurshid, X. Zou, W. Zhou, M. Caesar, and P. Brighten. VeriFlow: Verifying Network-wide Invariants in Real Time. In *Proceedings of the first Workshop on Hot Topics in Software-Defined Networks (HotSDN)*, 2012.

[112] H. Kim, J. Reich, A. Gupta, M. Shahbaz, N. Feamster, and R. Clark. Kinetic: Verifiable Dynamic Network Control. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation (USENIX NSDI)*, 2015.

[113] J. Kim, H. Choi, H. Namkung, W. Choi, B. Choi, H. Hong, Y. Kim, J. Lee, and D. Han. Enabling Automatic Protocol Behavior Analysis for Android Applications. In *Proceedings of the 12th ACM International Conference on Emerging Networking EXperiments and Technologies (ACM CONEXT)*, pages 281–295, New York, NY, USA, 2016.

[114] A. Kivity, Y. Kamay, D. Laor, U. Lublin, and A. Liguori. KVM: the Linux Virtual Machine Monitor. In *Proceedings of the Linux Symposium*, volume 1, pages 225–230.

[115] G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood. seL4: Formal Verification of an OS Kernel. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles (ACM SOSP)*, pages 207–220. ACM Press.

[116] J. Ko, S. Dawson-Haggerty, O. Gnawali, D. Culler, and A. Terzis. Evaluating the Performance of RPL and 6LoWPAN in TinyOS. In *Proceedings of the Workshop on Extending the Internet to Low Power and Lossy Networks (IP+SN)*, Chicago, IL, USA, April 2011.

[117] K. Kolyshkin. Virtualization in linux. *OpenVZ White Paper*, 3:39.

[118] M. La Polla, F. Martinelli, and D. Sgandurra. A Survey on Security for Mobile Devices. In *IEEE Communications Surveys & Tutorials*, 2012.

[119] A. Lahmadi, F. Beck, E. Finickel, and O. Festor. A Platform for the Analysis and Visualization of Network Flow Data of Android Environments. Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management (IFIP/IEEE IM).

[120] A. Lahmadi, A. Boeglin, and O. Festor. Efficient Distributed Monitoring in 6LoWPAN Networks. In *Proceedings of the IFIP/IEEE International Conference on Network and Service Management (IFIP/IEEE CNSM)*, Zürich, Switzerland, October 2013.

[121] A. Le, J. Loo, K. K. Chai, and M. Aiash. A Specification-Based IDS for Detecting Attacks on RPL-Based Network Topology. *Information*, 7(2), 2016.

[122] P. Levis, T. Clausen, J. Hui, O. Gnawali, and J. Ko. The Trickle Algorithm. RFC 6206 (Proposed Standard), Mar. 2011.

[123] C. Li, A. Raghunathan, and N. K. Jha. Secure Virtual Machine Execution under an Untrusted Management OS. In *Proceedings of the IEEE International Conference on Cloud Computing*, pages 172–179, 2010.

[124] D. Lie, C. A. Thekkath, and M. Horowitz. Implementing an Untrusted Operating System on Trusted Hardware. In *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles (ACM SOSP)*, pages 178–192. ACM Press.

[125] G. F. Lyon. *Nmap Network Scanning: The Official Nmap Project Guide to Network Discovery and Security Scanning*. Insecure, USA, 2009.

[126] A. Madhavapeddy, T. Leonard, M. Skjegstad, T. Gazagnaire, D. Sheets, D. J. Scott, R. Mortier, A. Chaudhry, B. Singh, J. Ludlam, and others. Jitsu: Just-In-Time Summoning of Unikernels. In *Proceedings of the 12th USENIX Symposium on Networked Systems Design and Implementation (USENIX NSDI)*, pages 559–573.

[127] A. Madhavapeddy, R. Mortier, C. Rotsos, D. Scott, B. Singh, T. Gazagnaire, S. Smith, S. Hand, and J. Crowcroft. Unikernels: Library Operating Systems for the Cloud. *SIGPLAN Notices*, 48(4):461–472, 2013.

[128] A. Madhavapeddy and D. J. Scott. Unikernels: Rise of the Virtual Library Operating System. *Queue*, 11(11):30:30–30:44, 2013.

[129] P. Manadhata. Game Theoretic Approaches to Attack Surface Shifting. In *Moving Target Defense: Application of Game Theory and Adversarial Modeling, Advances in Information Security*, 2013.

[130] A. Mayzaud. *Monitoring and Security for the RPL-based Internet of Things*. Phd Thesis, University of Lorraine, France, 2016.

[131] A. Mayzaud, R. Badonnel, and I. Chrisment. Monitoring and Security for the Internet of Things. In *Proceedings of the International Conference on Autonomous Infrastructure, Management and Security (IFIP AIMS)*, Barcelona, Spain, June 2013.

[132] A. Mayzaud, R. Badonnel, and I. Chrisment. A Taxonomy of Attacks in RPL-based Internet of Things. *International Journal of Network Security*, 18(3):459 – 473,, May 2016.

[133] A. Mayzaud, R. Badonnel, and I. Chrisment. Detecting Version Number Attacks in RPL-based Networks using a Distributed Monitoring Architecture. In *Proceedings of the IEEE/IFIP/In Assoc. with ACM SIGCOMM International Conference on Network and Service Management (IEEE/IFIP CNSM)*, Montreal, Canada, Oct. 2016.

[134] A. Mayzaud, A. Sehgal, R. Badonnel, I. Chrisment, and J. Schönwälder. A Study of RPL DODAG Version Attacks. In *Proceedings of the IFIP International Conference on Autonomous Infrastructure, Management and Security (AIMS)*, June 2014.

[135] A. Mayzaud, A. Sehgal, R. Badonnel, I. Chrisment, and J. Schönwälder. Mitigation of Topological Inconsistency Attacks in RPL-based Low-power Lossy Networks. *International Journal of Network Management*, 2015.

[136] A. Mayzaud, A. Sehgal, R. Badonnel, I. Chrisment, and J. Schönwälder. Using the RPL Protocol for Supporting Passive Monitoring in the Internet of Things. *In Proceedings of the IEEE/IFIP Network Operations and Management Symposium (IEEE/IFIP NOMS)*, Apr. 2016.

[137] M. McCormack, A. Vasudevan, G. Liu, S. Echeverría, K. O'Meara, G. Lewis, and V. Sekar. Towards an Architecture for Trusted Edge IoT Security Gateways. In *3rd USENIX Workshop on Hot Topics in Edge Computing (USENIX HotEdge)*. USENIX Association, June 2020.

[138] P. Mell and T. Grance. The NIST Definition of Cloud Computing. http://faculty.winthrop.edu/domanm/csci411/Handouts/NIST.pdf. Last visited on June 2021.

[139] MITRE Corporation. `http://www.mitre.org/`. Last visited on November 2021.

[140] Z. Movahedi, M. Ayari, R. Langar, and G. Pujolle. A Survey of Autonomic Network Architectures and Evaluation Criteria. *IEEE Communications Surveys & Tutorials*, PP:1–27, May 2011.

[141] L. Nelson, H. Sigurbjarnarson, K. Zhang, D. Johnson, J. Bornholt, E. Torlak, and X. Wang. Hyperkernel: Push-button verification of an OS kernel. In *Proceedings of the 26th Symposium on Operating Systems Principles (ACM SOSP)*, pages 252–269. ACM Press.

[142] NVD, National Vulnerability Database. `http://nvd.nist.gov/`. Last visited on November 2021.

[143] A. F. Ocampo, J. Gil-Herrera, P. H. Isolani, M. C. Neves, J. F. Botero, S. Latré, L. Zambenedetti, M. P. Barcellos, and L. P. Gaspary. Optimal Service Function Chain Composition in Network Functions Virtualization. In *Proceedings of the IFIP International Conference on Autonomous Infrastructure, Management and Security (IFIP AIMS)*, pages 62–76. Springer International Publishing, 2017.

[144] X. Ou, W. F. Boyer, and M. A. McQueen. A Scalable Approach to Attack Graph Generation. In *Proceedings of the 13th ACM Conference on Computer and Communications Security (ACM CCS)*, pages 336–345. ACM Press, 2006.

[145] X. Ou, S. Govindavajhala, and A. W. Appel. MulVAL: A Logic-based Network Security Analyzer. *on USENIX Security*, 2005.

[146] The OVAL Language. http://oval.mitre.org/. Last visited on November, 2020.

[147] C. Paetz. *Z-Wave Basics: Remote Control in Smart Homes*. CreateSpace Independent Publishing Platform, North Charleston, SC, USA, 2013.

[148] D. Palma and T. Spatzier. Topology and Orchestration Specification for Cloud Applications (TOSCA). *Organization for the Advancement of Structured Information Standards (OASIS), Tech. Rep*, 2013.

[149] M. Pattaranantakul, R. He, A. Meddahi, and Z. Zhang. SecMANO: Towards Network Functions Virtualization (NFV) Based Security MANagement and Orchestration. In *Proceedings of the IEEE Trustcom/BigDataSE/ISPA Conferences*, pages 598–605, Aug. 2016.

[150] M. Pattaranantakul, R. He, Z. Zhang, A. Meddahi, and P. Wang. Leveraging Network Functions Virtualization Orchestrators to Achieve Software-Defined Access Control in the Clouds. *IEEE Transactions on Dependable and Secure Computing*, 18(1):372–383, 2021.

[151] M. Pattaranantakul, Y. Tseng, R. He, Z. Zhang, and A. Meddahi. A First Step Towards Security Extension for NFV Orchestrator. In *Proceedings of the ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization (ACM SDN-NFVSec)*, pages 25–30. ACM.

[152] R. Patton. *Software Testing (2nd Edition)*. SAMS Publisher, 2005.

[153] B. D. Payne, M. Carbone, M. Sharif, and W. Lee. Lares: An Architecture for Secure Active Monitoring Using Virtualization. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 233–247, May 2008.

[154] M. Pearce, S. Zeadally, and R. Hunt. Virtualization: Issues, Security Threats, and Solutions. *ACM Computing Surveys*, 45(2):17:1–17:39, Mar. 2013.

[155] N. Petroulakis, K. Fysarakis, I. Askoxylakis, and G. Spanoudakis. Reactive Security for SDN/NFV-enabled Industrial Networks leveraging Service Function Chaining. *Transactions on Emerging Telecommunications Technologies*, Dec. 2017.

[156] D. Popa, N. Cam-Winget, and J. Hui. Applicability Statement for the Routing Protocol for Low Power and Lossy Networks (RPL) in AMI Networks. Internet-Draft Draft-ietf-roll-applicability-ami-13, Internet Engineering Task Force, May 2016. Work in Progress.

[157] G. J. Popek and R. P. Goldberg. Formal Requirements for Virtualizable Third Generation Architectures. *ACM Communications*, 17(7):412–421, July 1974.

[158] D. E. Porter, S. Boyd-Wickizer, J. Howell, R. Olinsky, and G. C. Hunt. Rethinking the library OS from the top down. 39(1):291–304.

[159] Puppet. http://www.puppetlabs.com/. Last visited on November 2021.

[160] K. N. Ramach, E. M. Belding-royer, and K. C. Almeroth. DAMON: A Distributed Architecture for Monitoring Multi-hop Mobile Networks. In *IEEE SECON*, Santa Clara, CA, USA, October 2004.

[161] S. Raza, L. Wallgren, and T. Voigt. SVELTE: Real-time intrusion detection in the Internet of Things. *Ad Hoc Networks*, 11(8):2661–2674, 2013.

[162] E. Rescorla and N. Modadugu. Datagram Transport Layer Security. RFC 4347 (Proposed Standard), Apr. 2006. Updated by RFC 5746.

[163] E. Rescorla and N. Modadugu. RFC 6347 (Proposed Standard): Datagram Transport Layer Security Version 1.2. https://datatracker.ietf.org/doc/html/rfc6347, Oct. 2015.

[164] R. Riley, X. Jiang, and D. Xu. Guest-Transparent Prevention of Kernel Rootkits with VMM-Based Memory Shadowing. In R. Lippmann, E. Kirda, and A. Trachtenberg, editors, *Proceedings of the International Symposium on Recent Advances in Intrusion Detection (RAID)*, pages 1–20. Springer Berlin Heidelberg, 2008.

[165] J. S. Robin and C. E. Irvine. Analysis of the Intel Pentium Ability to Support a Secure Virtual Machine Monitor. In *Proceedings of the 9th USENIX Security Symposium*, pages 129–144, 2000.

[166] D. Saha. Extending Logical Attack Graphs for Efficient Vulnerability Analysis. In *Proceedings of the 15th ACM Conference on Computer and Communications Security (ACM CCS)*, pages 63–74, New York, NY, USA, 2008. ACM.

[167] Sanjit A. Seshia. Verified Artificial Intelligence: A Runtime Verification Perspective, Keynote Talk at the International Conference on Runtime Verification (RV). 2019.

[168] N. Schnepf. *Orchestration et Verification de Fonctions de Securite pour des Environnements Intelligents*. Phd Thesis, University of Lorraine, France, 2019.

[169] N. Schnepf, R. Badonnel, A. Lahmadi, and S. Merz. *Automated Orchestration of Security Chains Driven by Process Learning*, chapter 12, pages 289–319. John Wiley and Sons, 2021.

[170] N. Schnepf, S. Merz, R. Badonnel, and A. Lahmadi. Automated Verification of Security Chains in Software-Defined Networks with Synaptic. In *Proceedings of the 3rd IEEE Conference on Network Softwarization (IEEE NetSoft)*, 2017.

[171] N. Schnepf, S. Merz, R. Badonnel, and A. Lahmadi. Rule-Based Synthesis of Chains of Security Functions for Software-Defined Networks. In *Proceedings of the 18th International Workshop on Automated Verification of Critical Systems (AVOCS)*, 2018.

[172] N. Schnepf, S. Merz, R. Badonnel, and A. Lahmadi. Towards Generation of SDN Policies for Protecting Android Environments based on Automata Learning. In *Proceedings of the 16th Network Operations and Management Symposium (IEEE/IFIP NOMS)*, 2018.

[173] N. Schnepf, S. Merz, R. Badonnel, and A. Lahmadi. Automated Factorization of Security Chains in Software-Defined Networks. In *Proceedings of the 16th IFIP/IEEE Symposium on Integrated Network and Service Management (IFIP/IEEE IM)*, 2019.

[174] S. Seeber, A. Sehgal, B. Stelte, G. D. Rodosek, and J. Schönwälder. Towards A Trust Computing Architecture for RPL in Cyber Physical Systems. In *IFIP/IEEE International Conference on Network and Service Management (IFIP/IEEE CNSM)*, Zürich, Switzerland, October 2013.

[175] A. Sehgal, A. Mayzaud, R. Badonnel, I. Chrisment, and J. Schönwälder. Addressing DODAG Inconsistency Attacks in RPL Networks. In *Proceedings of the Global Information Infrastructure Networking Symposium (GIIS)*, 2014.

[176] A. Sehgal, V. Perelman, S. Kuryla, and J. Schönwälder. Management of Resource Constrained Devices in the Internet of Things. *IEEE Communications Magazine*, 50(12):144–149, 2012.

[177] I. Sfyrakis and T. Großß. VirtusCap: Capability-Based Access Control for Unikernels. In *Proceedings of the IEEE International Conference on Cloud Engineering (IC2E)*, pages 226–237, Apr. 2017.

[178] H. Shacham, M. Page, B. Pfaff, E.-J. Goh, N. Modadugu, and D. Boneh. On the Effectiveness of Address-space Randomization. In *Proceedings of the 11th ACM Conference on Computer and Communications Security (ACM CCS)*, pages 298–307. ACM Press.

[179] A. Shameli-Sendi, Y. Jarraya, M. Pourzandi, and M. Cheriet. Efficient Provisioning of Security Service Function Chaining Using Network Security Defense Patterns. *IEEE Transactions on Services Computing*, Oct. 2016.

[180] SIDN Labs. Distributed Denial-of-Service (DDoS) Clearing House. `https://github.com/ddos-clearing-house`. Last visited on August 2021.

[181] S. Soltesz, H. Pötzl, M. E. Fiuczynski, A. Bavier, and L. Peterson. Container-based Operating System Virtualization: A Scalable, High-performance Alternative to Hypervisors. *SIGOPS Operating System Reviews*, 41(3):275–287.

[182] A. Sperotto. *Flow-based Intrusion Detection*. PhD thesis, University of Twente, 2010.

[183] U. Steinberg and B. Kauer. NOVA: A Microhypervisor-based Secure Virtualization Architecture. In *Proceedings of the European Conference on Computer Systems (EuroSys)*, pages 209–222. ACM Press.

[184] T. Sundqvist, M. H. Bhuyan, J. Forsman, and E. Elmroth. Boosted Ensemble Learning for Anomaly Detection in 5G RAN. In *Proceedings of the IFIP International Conference on Artificial Intelligence Applications and Innovations (IFIP AIAI)*.

[185] Apache Subversion. `http://subversion.apache.org/`. Last visited on November 2021.

[186] J. Szefer and R. B. Lee. A Case for Hardware Protection of Guest VMs from Compromised Hypervisors in Cloud Computing. In *Proceedings of the International Conference on Distributed Computing Systems*, pages 248–252, June 2011.

[187] U. Tupakula and V. Varadharajan. TVDSEC: Trusted Virtual Domain Security. In *Proceedings of the IEEE International Conference on Utility and Cloud Computing*, pages 57–64, Dec. 2011.

[188] The UniverSelf Project. `http://www.univerself-project.eu/`. Last visited on November 2021.

[189] C. Wagner, A. Dulaunoy, G. Wagener, and A. Iklody. MISP: The Design and Implementation of a Collaborative Threat Intelligence Sharing Platform. In *Proceedings of the ACM on Workshop on Information Sharing and Collaborative Security*, pages 49–56. ACM Press, 2016.

[190] A.-A. Walla. Live Updating in Unikernels. Master Thesis, University of Oslo, Norway, 2017.

[191] X. Wei. ProfileDroid : Multi-layer Profiling of Android Applications. In *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking (ACM MOBICOM)*, 2012.

[192] D. Williams and R. Koller. Unikernel Monitors: Extending Minimalism Outside of the Box. In *Proceedings of the 8th USENIX Workshop on Hot Topics in Cloud Computing (USENIX HotCloud)*. USENIX Association.

[193] T. Winter, P. Thubert, A. Brandt, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, J. Vasseur, and R. Alexander. RFC 6550: RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. https://datatracker.ietf.org/doc/html/rfc6550, 2012.

[194] X. Xu, J. Wan, W. Zhang, C. Tong, and C. Wu. PMSW: a passive monitoring system in wireless sensor networks. *International Journal of Network Management*, 21(4):300–325, 2011.

[195] G. Zhang, S. Ehlert, T. Magedanz, and D. Sisalem. Denial of Service Attack and Prevention on SIP VoIP Infrastructures using DNS Flooding. In *Proceedings of the International Conference on Principles, Systems and Applications of IP Telecommunications (IPTComm)*, pages 57–66, New York, NY, USA, 2007. ACM Press.

[196] F. Zhou, M. Goel, P. Desnoyers, and R. Sundaram. Scheduler Vulnerabilities and Coordinated Attacks in Cloud Computing. *Journal of Computer Security*, 21(4):533–559, 2013.

Glossary

6LoWPAN - IPv6 Low-power Wireless Personal Area Network.
AI - Artificial Intelligence.
AMI - Advanced Measurement Infrastructure.
AMPL - A Mathematical Programming Language.
AMQP - Advanced Message Queuing Protocol.
API - Application Programming Interface.
AT - Adaptive Threshold.
CNF - Conjunctive Normal Form.
CoAP - Constrained Application Protocol.
CoMI - CoAP Management Interface.
CPU - Central Processing Unit.
CSP - Cloud Service Provider.
CTL - Computation Tree Logic.
CVC - Cooperating Validity Checker.
CVE - Common Vulnerabilities and Exposures.
CVSS - Common Vulnerability Scoring System.
DAG - Directed Acyclic Graph.
DAMON - Distributed Architecture for MONitoring mobile network.
DAO - Destination Advertisement Object.
DAO-ACK - Destination Advertisement Object Acknowledgement.
DoS - Denial of Service.
DDoS - Distributed DoS.
DIO - DODAG Information Object.
DIS - DODAG Information Solicitation.
DNS - Domain Name System.
DODAG - Destination Oriented Directed Acyclic Graph.
DOVAL - Distributed OVAL.
DPI - Deep Packet Inspection.
DT - Dynamic Threshold.
EPMOSt - Energy-efficient Passive MOnitoring System.
FP - False Positive.
FPR - False Positive Rate.
FW - Firewall.
GSP - Global Security Policy.
IEEE - Institute of Electrical and Electronics Engineers.
IETF - Internet Engineering Task Force.
IDS - Intrusion Detection System.
ILP - Integer Linear Programming.

IOS - Cisco Internetworking Operating System.
IoT - Internet of Things.
IPS - Intrusion Prevention System.
JSON - JavaScript Object Notation.
KVM - Kernel-based Virtual Machine.
LLN - Low-power and Lossy Network.
MANET - Mobile Ad-hoc NETwork.
MIB - Management Information Base.
MISP - Malware Information Sharing Platform.
MMU - Memory Management Unit.
MOP - Mode Of Operation.
MP2P - Multipoint-to-Point.
MTD - Moving Target Defense.
NETCONF - NETwork CONFiguration protocol.
NFV - Network Function Vitualization.
OS - Operating System.
OVAL - Open Vulnerability and Assessment Language.
P2MP - Point-to-Multipoint.
P2P - Point-to-Point.
P4 - Programming Protocol-Independent Packet Processors.
PAP - Policy Administration Point.
PDP - Policy Decision Point.
PEP - Policy Enforcement Point.
PMSW - Passive Monitoring System for WSN.
RAM - Random-Access Memory.
RIP - Routing Information Protocol.
RoLL - Routing Over Low-power and Lossy networks.
RPL - Routing Protocol for Low-power and lossy networks.
SAT - Satisfiability.
SCAP - Security Content Automation Protocol.
SDK - Software Development Kit.
SDN - Software-Defined Networking.
SDSec - Software-Defined Security.
SFC - Service Function Chaining.
SIP - Session Initiation Protocol.
SMT - Satisfiability Modulo Theories.
SNMP - Simple Network Management Protocol.
SNTS - Sensor Networks Troubleshooting Suite.
SO - Security Orchestrator.
SVN - Subversion (Control Version System).
TLS - Transport Layer Security.
TLSP - Tenant-Level Security Policy.
TN - True Negative.
TOSCA - Topology and Orchestration Specification for Cloud Applications.
VM - Virtual Machine.
VMM - Virtual Machine Monitor.
VNF - Virtual Network Function.
WAN - Wide Area Network.

WSN - Wireless Sensor Network.
XACML - eXtensible Access Control Markup Language.
XCCDF - eXtensible Configuration Checklist Description Format.
XML - eXtensible Markup Language.
YOUNG - YOung Unikernel Generator.

Abstract

The Internet has become a great integration platform capable of efficiently interconnecting billions of entities, from simple sensors to large data centers. This platform provides access to multiple hardware and virtualized resources (servers, networking, storage, applications, connected objects) ranging from cloud computing to Internet-of-Things infrastructures. From these resources that may be hosted and distributed amongst different providers and tenants, the building and operation of complex and value-added networked systems is enabled. These systems are however exposed to a large variety of security attacks, that are also gaining in sophistication and coordination. In that context, the objective of my research work is to support security management for the cyberspace, with the elaboration of new monitoring and configuration solutions for these systems. A first axis of this work has focused on the investigation of smart monitoring methods capable to cope with low-resource networks. In particular, we have proposed a lightweight monitoring architecture for detecting security attacks in low-power and lossy networks, by exploiting different features provided by a routing protocol specifically developed for them. A second axis has concerned the assessment and remediation of vulnerabilities that may occur when changes are operated on system configurations. Using standardized vulnerability descriptions, we have designed and implemented dedicated strategies for improving the coverage and efficiency of vulnerability assessment activities based on versioning and probabilistic techniques, and for preventing the occurrence of new configuration vulnerabilities during remediation operations. A third axis has been dedicated to the automated configuration of virtualized resources to support security management. In particular, we have introduced a software-defined security approach for configuring cloud infrastructures, and have analyzed to what extent programmability facilities can contribute to their protection at the earliest stage, through the dynamic generation of specialized system images that are characterized by low attack surfaces. Complementarily, we have worked on building and verification techniques for supporting the orchestration of security chains, that are composed of virtualized network functions, such as firewalls or intrusion detection systems. Finally, several research perspectives on security automation are pointed out with respect to ensemble methods, composite services and verified artificial intelligence.

Keywords: Security Management, Cyberspace, Monitoring, Configuration.

Résumé

L'Internet est devenu une formidable plateforme d'intégration capable d'interconnecter efficacement des milliards d'entités, de simples capteurs à de grands centres de données. Cette plateforme fournit un accès à de multiples ressources physiques ou virtuelles, allant des infrastructures cloud à l'internet des objets. Il est possible de construire et d'opérer des systèmes complexes et à valeur ajoutée à partir de ces ressources, qui peuvent être déployées auprès de différents fournisseurs. Ces systèmes sont cependant exposés à une grande variété d'attaques qui sont de plus en plus sophistiquées. Dans ce contexte, l'objectif de mes travaux de recherche porte sur une meilleure gestion de la sécurité pour le cyberespace, avec l'élaboration de nouvelles solutions de monitorage et de configuration pour ces systèmes. Un premier axe de ce travail s'est focalisé sur l'investigation de méthodes de monitorage capables de répondre aux exigences de réseaux à faibles ressources. En particulier, nous avons proposé une architecture de surveillance adaptée à la détection d'attaques dans les réseaux à faible puissance et à fort taux de perte, en exploitant différentes fonctionnalités fournies par un protocole de routage spécifiquement développé pour ceux-ci. Un second axe a ensuite concerné la détection et le traitement des vulnérabilités pouvant survenir lorsque des changements sont opérés sur la configuration de tels systèmes. En s'appuyant sur des bases de descriptions de vulnérabilités, nous avons conçu et mis en œuvre différentes stratégies permettant d'améliorer la couverture et l'efficacité des activités de détection des vulnérabilités, et de prévenir l'occurrence de nouvelles vulnérabilités lors des activités de traitement. Un troisième axe fut consacré à la configuration automatique de ressources virtuelles pour la gestion de la sécurité. En particulier, nous avons introduit une approche de programmabilité de la sécurité pour les infrastructures cloud, et avons analysé dans quelle mesure celle-ci contribue à une protection au plus tôt des ressources, à travers la génération dynamique d'images systèmes spécialisées ayant une faible surface d'attaques. De façon complémentaire, nous avons travaillé sur des techniques de construction automatique et de vérification de chaînes de sécurité, qui sont composées de fonctions réseaux virtuelles telles que pare-feux ou systèmes de détection d'intrusion. Enfin, plusieurs perspectives de recherche relatives à la sécurité autonome sont mises en évidence concernant l'usage de méthodes ensemblistes, la composition de services, et la vérification de techniques d'intelligence artificielle.

Mots-clés: Gestion de la sécurité, Cyberespace, Monitorage, Configuration.