Security Automation in next-generation Networks and Cloud environments

Security Automation in next-generation Networks and Cloud environments

Francesco Pizzato, Daniele Bringhenti, Riccardo Sisto, Fulvio Valenza
Dip. Automatica e Informatica, Politecnico di Torino, Torino, Italy, Emails: {first.last}@polito.it

***Abstract*—In the next generation networks and cloud systems, administrators should only need to define their intentions through simple high-level intents, leaving the system to autonomously implement them in the best way possible. The adoption of automation enables the possibility to create reactive systems that can reconfigure themselves in response to unpredictable events, such as network attacks. Nowadays, such solutions are far from being achieved. The enforcement of security requirements continues to heavily rely on manual efforts and tools requiring non-negligible expertise to be used. This results in frequent misconfiguration errors or the complete absence of default security measures due to their high implementation complexity. This paper introduces the research that will be carried out within my Ph.D. program, focusing on network security automation. The objective is to bridge existing gaps in the literature, on one side developing novel automated and intent-based approaches for security enforcement in cloud environments, ensuring formal correctness and optimization, and on the other side researching new solutions for the design of security reaction mechanisms for modern networks in response to network attacks.**

***Index Terms*—security automation, optimized reconfiguration, cloud security**

I. Introduction and motivation

Cloud Computing, as defined by NIST, represents a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. Today, seamless management and automation have been extensively achieved only for the handling of computing resources. For instance, modern cloud platforms implement automated solutions for resource scaling, workload management, and resource provisioning. However, the configuration of security services in the cloud continues to strongly depend on human effort and expertise for proper implementation.

Furthermore, modern cloud-native applications should be designed to leverage the advantages of the cloud computing model, emphasizing scalability, elasticity, and automation. Best practices, such as "12-Factors" [1], advocate the usage of declarative models and automation to minimize development time and cost. Treating cloud security with the same best practices, i.e., automation and declarative approach, could be crucial to address most of today's challenges, mitigating the risk of security misconfigurations, and reducing the burden on security professionals. As reported by different sources, misconfiguration of cloud services and resources is the second most common factor involved in attacks with a frequency of 42.6% [1], and it is the number one threat requiring the larger portion of security experts' effort in the day-to-day tasks for 59% of the interviewed professionals [2].

In the future generation of cloud environments, security services should be managed in the same way as other cloud resources. In particular, the user should be able to define, through high-level declarative intents, the desired security posture for the system, and the interactions between him and the underlying service provider should be minimal. Optimization and automation are crucial to implement this vision and to solve many of the today's challenges. Optimization, a feature already considered in cloud workload orchestration [3], can further enhance automated approaches as it allows for the reduction of resource utilization by selecting the optimal placement and configuration of the needed security functions. Also, the usage of formal methods, could aid automated approaches by providing correctness assurance on the computed configuration and the enforcement of the requested security requirements.

In light of these considerations, the main goal of my Ph.D. is to advance the state of the art of security automation and optimization in next-generation networks and cloud-based systems. In particular, I will focus on the design of novel approaches for the automated implementation of cloud security services, such as access control, isolation policies, and key security principles (e.g., least privilege and zero trust). Despite there being some tools available today offering partial support for these activities, they still have some limitations and impose a substantial manual burden on users, who must write several policies and configurations manually, with the risk of introducing errors that undermine security. Therefore, a complementary research path will be addressing the design of enhanced re-configuration approaches and reactive defense systems. Accounting for the high dynamism of modern networks and cloud environments, being able to rapidly react and re-configure the security posture of the system is another aspect with high potential impact.

The remainder of this paper is organized as follows. Section II briefly introduces some relevant concepts constituting the foundations of my research. Section III outlines the research direction and a brief presentation of prior and current works. Finally, Section IV contains the conclusions and the next steps in my research plan.

[1] "The Twelve-Factor App", Wiggins Adam, 2017, https://12factor.net/

II. Background

A. Network Security Automation

The size and complexity of modern networks have steeply grown in the last few years, presenting new challenges for the implementation of network security. Additionally, the traditional way of securing computer networks involves manual trial-and-error approaches. Whenever an attack is detected, the security manager has to manually update the configuration. This results in poor scalability of the traditional methods and highlights the human inability to manage complex systems, as it has been reported by several sources, for instance, the last report from Verizon [4] reported that 74% of the analyzed breaches involved a human element among the identified causes. To address this problem, new approaches have been proposed to implement network security automation, i.e., solutions capable of automatically designing and configuring the necessary network security functions (e.g., firewalls, VPNs). These innovative solutions allow for the reduction of configuration time, management of complex systems, and minimization of human intervention, thereby mitigating the risk of misconfigurations.

Several such solutions have been published in the literature. A comprehensive analysis about the state of the art of automatic security orchestration is presented in [5]. According to the authors, the different solutions can be classified based on some key features, such as the kind of network security functions that can be configured, the capability of providing a formal assurance about the solution's correctness, or the support of optimization criteria and which approach has been selected to achieve them. The most relevant ones are ConfigSynth [6] and VEREFOO [7], [8], as they are both supporting a complete set of features, that is automation, optimization and formal correctness assurance, for the automatic configuration of distributed packet filter firewalls.

B. Policy-Based Management

The effectiveness of network security automation is significantly enhanced when it is integrated with *Policy-Based Management* (PBM) [9], i.e., processing the management operations using policies. This combination allows to achieve a comprehensive strategy for the declarative definition and automated enforcement of security policies. The process of PBM unfolds in three key phases. The initial one involves the formulation of user-specified policies and their automated analysis for anomaly detection, this phase is referred to as *policy analysis*. Subsequently, the second phase is *policy refinement*, consisting of refining policies into low-level configuration rules specific to the actual implementation. The final step consists of the verification of the produced configuration, to ensure compliance with the original policies defined by the user. This is denoted as *policy verification*. An additional step which completes the policy-based workflow is *security mitigation*. Once the network is configured and policies are correctly enforced, the automated approach collaborates with monitoring agents (e.g., intrusion detection systems). Whenever a cyber attack is identified, it triggers an automated mitigation process involving the generation of new security policies and the subsequent reconfiguration of the involved network security functions. It's important to note that, in our research, we emphasize the design of solutions that are formally correct by construction. In this way, one of the three main phases of PBM, policy verification, could be ignored since the solution is implicitly verified due its guaranteed correctness.

C. Policy Languages

Based on the analyses conducted by Basile et al. [10] and Hermosilla et al. [11], a comprehensive representation of all security configurations can be achieved with three classes of policy languages. *High-level Policy Languages* (HPLs) enable users to articulate policies in a user-friendly notation, thereby enhancing readability and comprehension. *Medium-level Policy Languages* (MPLs) offer a structured, implementation-independent representation of policies, utilizing a combination of conditions (events triggering the policy) and actions (operations executed when conditions are met). *Low-Level Configurations* involve languages tailored to the specific requirements of the adopted implementation for the network functions. In this context, the refinement process of PBM is responsible for transitioning from higher to lower classes of policy languages. By combining all these policy languages, users can define policies in an user-friendly language, and the automated tool can convert it into a low-level language suitable for the specific implementation of the network functions.

III. Research Direction

Considering all the reasons that have been discussed so far, the research activity within my Ph.D. program will examine security automation for the next-generation networks and cloud systems. This objective has been partially investigated in this initial phase, with a preliminary analysis of the literature about cloud security and automation, as presented in III-A, and with the design of an efficient reconfiguration algorithm able to quickly re-compute the security configuration of a network in the event of an attack, representing an initial step to build a complete reactive approach for the automated reconfiguration of network security functions. This is presented in III-B.

A. Cloud Security Automation

A first research area is the design of automated and optimal approaches for security automation within cloud environments. This idea stems from what has been successfully achieved in network security automation, as a similar approach could be considered in the context of cloud-based systems, characterized by improved dynamism and flexibility. In particular, this requires the design of formal models capable of exhaustively representing cloud environments, and the configuration of cloud security services. Another key component is the definition of a suitable language for user-defined security intents and the design a refinement process able to derive low-level configurations starting from the defined intents. The final goal is an approach able to automatically generate an optimal and

correct security posture through the combination of different cloud security services (e.g., RBAC rules, Network Policy). The usage of formal methods is necessary as it provides formal correctness assurance about the computed configuration. The research will mainly focuses on Kubernetes, being the de-facto standard orchestration platform for cloud workloads. Even though it supports different security features, navigating through them and formulating a comprehensive security strategy can be a daunting task, often leading to misconfigurations that significantly undermine the desired security posture of the system. For these reasons, we are convinced that the application of network security automation within the scope of Kubernetes represents a novelty with respect to the current state of the art and has a high potential impact due to the wide adoption of the technology and its challenging aspects concerning security management.

An initial review of the scientific literature identified the lack of studies in this area, representing an interesting niche not yet explored. To the best of our knowledge, few publications are considering Kubernetes, and only some of them adopts automation and formal models. None of them is proposing a solution for automatic security configuration. A relevant number of papers [12]–[14], authored by researchers within AWS Science, are proposing the usage of formal methods to automate verification and compliance of security policies within the AWS cloud platform. For instance, [14] proposes an approach for network reachability verification using theorem provers and formal models describing different AWS networking components and their configuration. Instead, [13] covers the design of an analysis tool for the automatic verification of access control rules, defined with an AWS-specific policy language, which is based on the design of formal models and the resolution through different SMT solvers. [12] presents an approach for the execution of a vulnerability analysis on the system before it is deployed. Kubernetes supports the usage of Infrastructure-as-a-Code files to descriptively define and deploy a complete system. Through the analysis of these files, the presented approach can extract a model representing the cloud system, which is subsequently enriched by analyzing also the configuration files. This is an interesting area of research that has been investigated in similar studies, such as [15], [16]. In general, the research effort towards the application and verification of security services in Kubernetes is not enough. Moreover, no current solution includes a completely automated approach capable of configuring the security functions in the cloud starting from a set of high-level policies.

Ongoing work in this area regards the design of an automated intent-based solution for the enforcement of a network security perimeter in the case of a multi-cluster and multi-tenant environment. This idea was developed as a contribution to the FLUIDOS project, which scope is to create a meta-OS based on the liquid computing model, i.e., an evolution of cloud computing in which different tenants share a continuum pool of resources that could be borrowed or lent in a ”liquid“ way. In this context, the concept of a fixed physical boundary is substituted by ephemeral and evolving

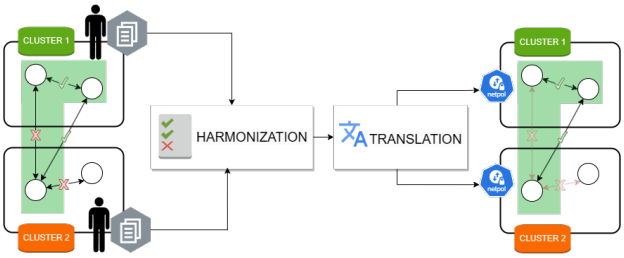

Fig. 1: Schema of the intent-based network isolation workflow.

virtual borders, requiring the design of novel and innovative approaches to enforce network isolation. In greater detail, the proposed solution allows the formulation and enforcement of finely-grained isolation policies for the network, to ease the implementation of common security patterns such as zero trust and least privilege. A high-level description of the workflow is represented in Fig. 1. The importance of the work is highlighted by recent surveys about the security of Kubernetes systems, for instance [17] reported that among all the analyzed clusters, only 9% of them have namespaces with configured Network Policy, which is a recommended configuration to prevent lateral movements of attackers inside the cluster.

B. Reactive security

A second aspect of my research considers the design of a policy-based reaction system to efficiently make use of the dynamism of modern networks and policy-based workflows. The initial analysis of the state of the art highlighted that it is still missing a fully automated and fast enough reaction solution. Most of the existing approaches for network security automation do not support the reconfiguration scenario and so require the re-computation of both allocation and configuration of each security functions in the entire network from scratch every time the security policies are updated. This results in time-consuming reconfigurations that are not compatible with the requirement of fast reaction to attacks. Moreover, also the way in which the update is carried out could be subject to unsafe states and so it can potentially benefit from optimizations, as described in [18].

In this area, we already developed a new fast reconfiguration algorithm for distributed firewalls in the context of virtualized networks, starting from the state-of-the-art VEREFOO (Verified REFinement and Optimized Orchestration) approach [7], [8]. The preliminary results of this work have been included in a paper that will be presented at the IEEE/IFIP NOMS 2024 conference [19]. Another research work, which is currently under development, logically follows the previous one and focuses on the design of an automated mechanism to connect the outputs of intrusion detection algorithms (i.e., the alarms produced by the monitoring and detection system) to the automatic policy-based security reconfiguration process. To create this connection, the main challenge is the design of a policy extraction engine capable of extracting crucial information from the alerts produced by monitoring and detection agents and then producing a set of

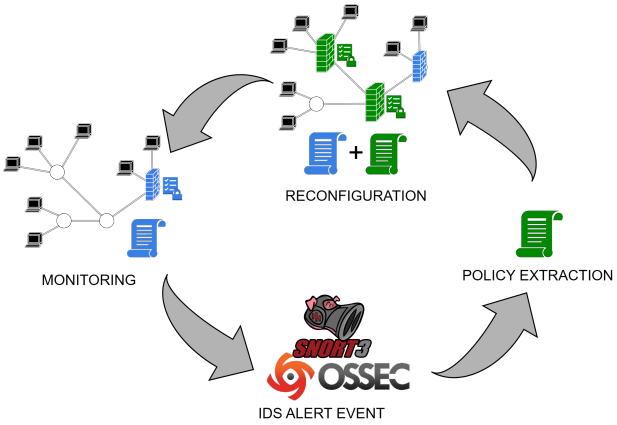

Fig. 2: general schema of the reactive approach

security intents that correctly resolve the detected issue. An initial prototype of the engine has been developed with two of the most widely adopted IDS systems, i.e., Snort3 and OSSEC 3.7, and the evaluation is currently in progress. This allows us to achieve a comprehensive approach to automatically reconfigure a network in the event of an attack without needing human inputs, as represented in Fig. 2.

C. Research areas for improvements

To conclude, the two research areas that I have identified of interest for my Ph.D. research program are the following:

1) Design of an approach for the automatic and optimal enforcement of access control in Kubernetes using an intent-based workflow. Among the many security services, I think that access control could be interesting because it has a very important role in the security hardening of a system, and because it is an area where similar approaches have already been employed in traditional and virtualized networks.
2) Further research towards the design of the reactive-security framework. This area could be expanded along two different directions. Firstly, investigate potential ways to integrate a similar framework within Kubernetes. Secondly, focus on enhancing a critical aspect of this kind of systems, the policy extraction engine. The process of generating security policies from security alerts is a key element that is often carried out in a simplistic manner, and there is potential for improvement to optimize this process.

IV. Conclusion

In this paper, I have discussed the challenges and open problems in security automation for the next generation of cloud environments. I have also analyzed the direction of my Ph.D. program, which is the design of automated approaches for the optimized and formal configuration of cloud security, and the parallel interest in developing a reactive framework to automate and improve the re-configuration of a security system in the event of an attack. My previous and ongoing work has also been presented in Section III as it represents a reasonable base for the progression towards the proposed research areas. In the future, I want to delve deeper into the topics and continue to work on both areas, more specifically in the context of cloud environments and Kubernetes-based systems.

References

[1] SANS, "2022 Cloud Security Survey," 2022, Available: https://www.sans.org/white-papers/sans-2022-cloud-security-survey/, Visited: 2024-01-15.
[2] Fortinet, "2023 Cloud Security Report," 2023, Available: https://global.fortinet.com/emea-lp-en-2023cloudsecurityreport, Visited: 2024-01-15.
[3] M. Masdari, S. S. Nabavi, and V. Ahmadi, "An overview of virtual machine placement schemes in cloud computing," *J. Netw. Comput. Appl.*, vol. 66, pp. 106–127, 2016.
[4] Verizon, "2023 Data Breach Investigations Report," 2023, Available: https://www.verizon.com/business/resources/reports/2023-data-breach-investigations-report-dbir.pdf, Visited: 2024-01-15.
[5] D. Bringhenti, G. Marchetto, R. Sisto, and F. Valenza, "Automation for network security configuration: State of the art and research trends," *ACM Comput. Surv.*, vol. 56, no. 3, pp. 57:1–57:37, 2024.
[6] M. A. Rahman and E. Al-Shaer, "Automated synthesis of distributed network access controls: A formal framework with refinement," *IEEE Trans. Parallel Distributed Syst.*, vol. 28, no. 2, pp. 416–430, 2017.
[7] D. Bringhenti, G. Marchetto, R. Sisto, F. Valenza, and J. Yusupov, "Automated optimal firewall orchestration and configuration in virtualized networks," in *IEEE/IFIP Network Operations and Management Symposium (NOMS), Budapest, Hungary, April 20-24, 2020*. IEEE, 2020, pp. 1–7.
[8] ——, "Automated firewall configuration in virtual networks," *IEEE Trans. Dependable Secur. Comput.*, vol. 20, no. 2, pp. 1559–1576, 2023.
[9] R. Boutaba and I. Aib, "Policy-based management: A historical perspective," *J. Netw. Syst. Manag.*, vol. 15, no. 4, pp. 447–480, 2007.
[10] C. Basile, F. Valenza, A. Lioy, D. R. López, and A. P. Perales, "Adding support for automatic enforcement of security policies in NFV networks," *IEEE/ACM Trans. Netw.*, vol. 27, no. 2, pp. 707–720, 2019.
[11] A. Hermosilla, A. M. Zarca, J. B. Bernabé, J. O. Murillo, and A. F. Skarmeta, "Security orchestration and enforcement in nfv/sdn-aware UAV deployments," *IEEE Access*, vol. 8, pp. 131 779–131 795, 2020.
[12] C. Cauli, M. Li, N. Piterman, and O. Tkachuk, "Pre-deployment security assessment for cloud services through semantic reasoning," in *Proc. of Computer Aided Verification (CAV) - 33rd International Conference, Virtual Event, July 20-23, 2021*, ser. Lecture Notes in Computer Science, vol. 12759. Springer, 2021, pp. 767–780.
[13] J. Backes *et al.*, "Semantic- based automated reasoning for aws access policies using smt," in *Formal Methods in Computer Aided Design (FMCAD)*, 2018, pp. 1–9.
[14] ——, "Reachability analysis for aws-based networks," in *Proc. of Computer Aided Verification (CAV) - 31st International Conference, New York City, NY, USA, July 15-18, 2019*, ser. Lecture Notes in Computer Science, vol. 11562. Springer, 2019, pp. 231–241.
[15] A. Blaise and F. Rebecchi, "Stay at the helm: secure kubernetes deployments via graph generation and attack reconstruction," in *IEEE 15th Int. Conf. on Cloud Computing (CLOUD), Barcelona, Spain, July 10-16, 2022*. IEEE, 2022, pp. 59–69.
[16] F. Minna, F. Massacci, and K. Tuma, "Towards a security stress-test for cloud configurations," in *IEEE 15th Int. Conf. on Cloud Computing (CLOUD), Barcelona, Spain, July 10-16, 2022*. IEEE, 2022, pp. 191–196.
[17] Wiz, "The 2023 Kubernetes Security Report," 2023, Available: https://www.wiz.io/lp/the-2023-kubernetes-security-report, Visited: 2024-01-15.
[18] D. Bringhenti and F. Valenza, "Optimizing distributed firewall reconfiguration transients," *Comput. Networks*, vol. 215, p. 109183, 2022.
[19] F. Pizzato, D. Bringhenti, R. Sisto, and F. Valenza, "Automatic and optimized firewall reconfiguration," to appear in NOMS 2024 - IEEE/IFIP Network Operations and Management Symposium, Seoul, Korea, May 6-10, 2024.