Inside Job: Defending Kubernetes Clusters Against Network Misconfigurations

JACOPO BUFALINO, CNAM, France and Aalto University, Finland
JOSE LUIS MARTIN-NAVARRO, Universitat Politècnica de València, Spain and Aalto University, Finland
MARIO DI FRANCESCO, Aalto University, Finland
TUOMAS AURA, Aalto University, Finland

Kubernetes has emerged as the de facto standard for container orchestration. Unfortunately, its increasing popularity has also made it an attractive target for malicious actors. Despite extensive research on securing Kubernetes, little attention has been paid to the impact of network configuration on the security of application deployments. This paper addresses this gap by conducting a comprehensive analysis of network misconfigurations in a Kubernetes cluster with specific reference to lateral movement. Accordingly, we carried out an extensive evaluation of 287 open-source applications belonging to six different organizations, ranging from IT companies and public entities to non-profits. As a result, we identified 634 misconfigurations, well beyond what could be found by solutions in the state of the art. We responsibly disclosed our findings to the concerned organizations and engaged in a discussion to assess their severity. As of now, misconfigurations affecting more than thirty applications have been fixed with the mitigations we proposed.

CCS Concepts: • **Security and privacy** → **Distributed systems security**; **Network security**.

Additional Key Words and Phrases: Kubernetes, lateral movement, misconfigurations, deployments, Helm

ACM Reference Format:
Jacopo Bufalino, Jose Luis Martin-Navarro, Mario Di Francesco, and Tuomas Aura. 2025. Inside Job: Defending Kubernetes Clusters Against Network Misconfigurations. 1, 1 (June 2025), 26 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

The security of cloud applications is a primary concern [36]. Protecting cloud environments from malicious actors presents unique challenges that are not only technology-dependent, but rather originate from the intrinsic characteristics of cloud-native applications – those designed from the ground-up to take full advantage from the cloud paradigm [68]. In fact, modern cloud-based applications comprise a large number of containerized microservices, that is, loosely coupled software components that leverage application programming interfaces (APIs) over a network [60]. Kubernetes [24] has become the de-facto standard for container orchestration in this context. Unfortunately, the security of Kubernetes applications is particularly worrying: a report recently released by Red Hat and based on a survey of 600 IT professionals revealed that 89% of them experienced at least one security incident, even resulting in revenue or customer loss [67].

Authors' Contact Information: Jacopo Bufalino, CNAM, Paris, France and Aalto University, Espoo, Finland, jacopo.bufalino@lecnam.net; Jose Luis Martin-Navarro, Universitat Politècnica de València, València, Spain and Aalto University, Espoo, Finland, jose.martinnavarro@aalto.fi; Mario Di Francesco, Aalto University, Espoo, Finland, mario.di.francesco@aalto.fi; Tuomas Aura, Aalto University, Espoo, Finland, tuomas.aura@aalto.fi.

ACM XXXX-XXXX/2025/6-ART
https://doi.org/10.1145/nnnnnnn.nnnnnnn

```yaml
containers:
  - name: flink
    image: bitnami/flink
    ports:
      - containerPort: 6121
      - containerPort: 6123
      - containerPort: 8081
```

(a)

```
# netstat -a
Active Internet connections (servers and established)
Proto  Recv-Q  Send-Q  Local Address  Foreign Address  State
tcp         0       0  0.0.0.0:6123   0.0.0.0:*        LISTEN
tcp         0       0  0.0.0.0:8081   0.0.0.0:*        LISTEN
tcp         0       0  0.0.0.0:43271  0.0.0.0:*        LISTEN
```

(b)

Fig. 1. Configuration mismatch involving the Apache Flink application, a widely used data-flow framework for stream processing, as packaged by Bitnami. (a) The corresponding Kubernetes configuration (i.e., YAML file) specifies 6,121, 6,123 and 8,081 as the container ports; (b) the application is actually not listening to port 6,121 but uses an ephemeral port (i.e., 43,271), as reported by the `netstat` tool. These kinds of mismatches combined with application misbehavior may lead to different types of security issues (see Sections 2.1 and 3.3).

Specifically, misconfigurations have been recognized as primary causes of attacks to Kubernetes [1, 31, 67, 78], also affecting Fortune 500 companies [6]. Misconfigurations occur when the deployment of Kubernetes application is poorly setup (Figure 1), typically by administrators instead of the actual developers of these applications. The danger of such misconfigurations is amplified by the increasing trend of reusing third-party resources and configurations [19]. Furthermore, many related issues arise from the cloud networking paradigm, which is substantially different from legacy corporate networks – those using physical cables and on-premise hardware, including firewalls and switches, for segmentation purposes [74]. The Kubernetes model in particular is based on a flat network to simplify the transition of legacy applications to the cloud, however, its security implications are not fully understood [55]. As a consequence, it is challenging to correctly configure all the components of a cloud application, which remains a tedious and error-prone process [65].

Research on Kubernetes security has largely focused on hardening [70], risk assessment [13], and tools for runtime analysis [56]. Instead, studies on security misconfigurations have primarily targeted the cloud infrastructure [11, 18, 64] and container runtimes [40]. In fact, only a few works have explicitly addressed misconfigurations in the context of Kubernetes security [34, 65, 67, 79]. The survey-based report in [67] established misconfigurations as one of the major factors behind security threats in Kubernetes, but did not further elaborate on their characteristics. The empirical study in [65] targeted security misconfigurations in open-source Kubernetes applications, whereas the works of [34, 79] focused on excessive permissions granted to third-party components. However, the existing literature has paid little attention to security issues originated by network misconfigurations, in particular, those affecting actual cluster deployments (Section 7).

This work explicitly characterizes network misconfigurations affecting Kubernetes applications in the context of lateral movement (Section 2). Specifically, we rigorously examined research papers and industry artifacts as a basis for independent research that resulted in a comprehensive list of network misconfigurations (Section 3). Then, we carried out an analysis of 287 open-source Kubernetes applications belonging to six different organizations, ranging from IT companies and public entities to non-profits. The analysis revealed a total of 634 misconfigurations, most of which cannot be detected by solutions in the state of the art (Section 4). Finally, we responsibly disclosed our findings to the respective organizations and assessed their severity. As of now, misconfigurations affecting more than thirty applications have been fixed with the mitigations we proposed (Section 5).

In summary, the contributions of this paper are the following.

- **Novel network misconfigurations leading to security issues**. Our analysis of Kubernetes cluster-internal networking unveiled different types of misconfigurations, well beyond what

was known in the existing literature. We also show how these misconfigurations, together with application (mis)behavior, can lead to diverse security issues.
- **An evaluation of these misconfigurations in real-world application deployments**. We considered 287 open-source applications belonging to different organizations as deployed into a Kubernetes cluster. Our evaluation allowed to detect 634 network misconfigurations. We also compare the effectiveness of our solution against the state of the art.
- **Disclosure and community engagement**. We disclosed the misconfigurations to the involved organizations and explained how to mitigate them. We also carried out a security assessment and characterized the reliability of our results through the received feedback.

2 Background

This section provides a few examples of attacks we discovered to motivate the study of network misconfigurations. It then gives an overview of Kubernetes and application deployment with Helm.

2.1 Motivating Examples: Security Issues Resulting from Network Misconfigurations

We describe next two proof of concept attacks that are enabled by network misconfigurations.

2.1.1 Broken Control Plane: Concourse. Concourse [29] is a CI/CD software part of the CNCF landscape [21], consisting of a master web node and a number of different workers that are responsible for running builds and carrying out different types of resource checks. The Concourse application opens several dynamic ports in the ephemeral port range reserved by the host operating system (e.g., 32,768–60,999). These ports in the Concourse web node are endpoints of reverse SSH tunnels to worker nodes. The tunnels are command and control channels to the workers which should only be available at the loopback adapter of the web node. Instead, the ports are accessible from the cluster network due to the default Kubernetes behavior. Thus, any pod in the same cluster as the web node is able to send commands to a given worker, including other workers. As a result, we were able to deploy arbitrary containers, download container images, and edit running jobs.

2.1.2 Service Impersonation: Thanos. Thanos [77] is a set of components realizing high availability and long-term storage for the widely-used Prometheus monitoring system [27]. The related setup includes two compute units: `thanos-query-frontend`, responsible for external communication with users; and `thanos-query`, handling internal query processing. Both of them run different services that are associated with a single label, namely, `thanos-query-frontend`. As a consequence, an attacker inside either compute unit could exploit such a label to impersonate the service by simply listening to the appropriate ports. In fact, the load balancer or service targeting the resources bound to that label would send requests to the malicious pods in addition to (instead of) the legitimate ones, according to the applicable scheduling policy (i.e., depending on the specific proxy mode set in Kubernetes). Therefore, the attacker could fully impersonate the service, resulting in a denial of service (when the malicious pod does not respond) or even economic damage (by triggering unnecessary scale-up of cluster resources).

2.2 Kubernetes and Container Networking

Kubernetes [24] is the most widely used open-source software for container orchestration: it allows to deploy, scale and manage applications defined as interconnected microservices (Figure 2a). Kubernetes relies on a set of *nodes* (either bare-metal servers or virtual machines) together forming a *cluster*. One node in the cluster has a special role: such a *control-plane node* includes an API server (to interact with the orchestrator), a controller manager (to enforce desired attributes), and a scheduler (to allocate computing resources to nodes). The rest of the nodes, instead, run the actual applications (i.e., workloads). All resources in a cluster are specified by configuration files in the

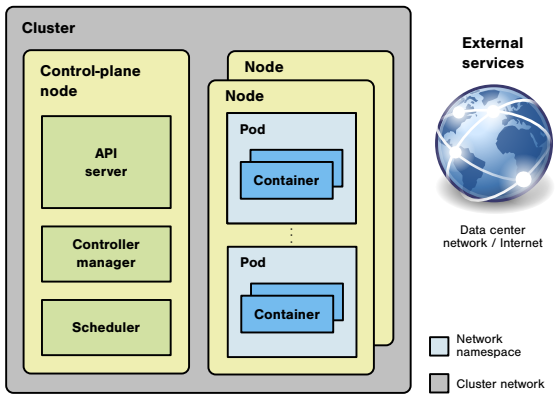

(a)

```yaml
# Helm manifest (values.yaml)
primary:
  service:
    type: ClusterIP
    ports:
      mysql: 3306
```

```yaml
# Helm template
kind: Service
metadata:
  labels:
    app.kubernetes.io/part-of: mysql
  ports:
    - name: mysql
      port: {{ .Values.primary.service.ports.mysql }}
      protocol: TCP
```

(b)

Fig. 2. (a) The main components in a cluster running microservices with Kubernetes. (b) Sample fragments of a Helm manifest (top) and a Helm template (bottom). The manifest defines values that can be referenced by the corresponding keys. The template in the example is a service and includes a directive (enclosed in double curly brackets) that sets the `port` key to the value of `mysql`. Such a value is specified as `.Values.primary.service.ports.mysql`, according to the structure of the manifest.

YAML format; computing resources are represented by software containers. Kubernetes objects are identified by their name and a list of key-value pairs called *labels* [45]. Kubernetes manages application lifecycle through a container *runtime* (e.g., docker or podman) installed on the node to retrieve container images and create their execution environment. In particular, Kubernetes employs the abstraction of *pod* to manage logically-related containers as an individual entity.

Pods communicate with each other by exchanging messages over a network. The Kubernetes networking model is grounded on a few key abstractions [55]. Containers in the same pod are bridged to a shared Linux network namespace which allows communication over the same *local network* (i.e., the so-called localhost). The address space in the cluster is flat, namely, pods can communicate with all other pods irrespective from the node where they are running. The *service* resource acts as an abstract representation of a logical set of pods for load-balancing and service discovery, for instance. Pods are assigned to each service by means of their labels. Services can be cluster-internal or external, depending on the specific use case. Cluster-internal services are defined by the `ClusterIP` type and are assigned a unique IP address in the cluster network. As a special case, a *headless* service [44] allows to reach pods through the cluster's Domain Name System (DNS). Pods in the cluster can communicate with all services by default.

Kubernetes also supports network policies (i.e., through the `NetworkPolicy` object) to restrict access to pods and services. However, Kubernetes itself is not concerned about how actual IP addresses are managed and how network policies are enforced – these are instead handled by external components called Container Network Interface (CNI) plugins.

2.3 Managing Kubernetes Applications

In the context of Kubernetes, applications are just a collection of resources specified in YAML files, such as the sample in Figure 1a. Managing such files is tedious and error-prone, especially when the same resources need to be tailored for environments with varying characteristics, such as for development or testing purposes. For this reason, applications are usually specified through reusable Kubernetes configurations (generally called templates or blueprints) with specialized

tools. Among them, Helm [22] is a widely used software to package and manage applications for Kubernetes clusters; a Helm *chart* is a collection of files that describes applications as Kubernetes resources, including configuration values and dependencies [35]. Specifically, a chart includes a manifest (specifying parameter values) and a template (describing a Kubernetes resource), as illustrated in Figure 2b. The templates are rendered into Kubernetes objects as YAML files by setting the parameters to the corresponding values. Helm charts are also composable, allowing to reuse other applications as dependencies – many of them are publicly available and are specifically meant for sharing [20]. The terms (Kubernetes) application and (Helm) chart are used interchangeably in the rest of the paper.

3 Network Misconfigurations

This section describes the network misconfigurations that can lead to security issues in Kubernetes. It first introduces the reference threat model, then explains the methodology employed to identify these misconfigurations. The section concludes with a detailed discussion of individual misconfigurations, along with an explanation of their causes and possible mitigations.

3.1 Threat Model

We explicitly refer to the threat matrix for Kubernetes by Microsoft [53], a widely-used resource in the context of container orchestration security, based on the MITRE ATT&CK framework [57]. We restrict our attention to the lateral movement tactic in the threat matrix, particularly, the "*cluster-internal networking*" technique therein. Specifically, we aim to identify and characterize misconfigurations possibly causing security threats in such a context for applications as they are deployed in a Kubernetes cluster. For this purpose, we assume that the attacker controls one container in a pod, which is aligned with previous work on Kubernetes security [34]. Moreover, our work assumes that the cluster has been hardened according to security best practices [10], namely, that it is not possible for the attacker to trivially take control of the whole cluster or a node. Finally, the container in control of the attacker has legitimate access to the cluster network but no other privileges that can be misused (e.g., root access or Kubernetes API access).

3.2 Methodology

We carried out a systematic review by collecting sources on network security for Kubernetes from academia and industry similar to the methodology in [2, 65], as detailed next. In particular, we determined misconfigurations based on the corresponding definition by the U.S. National Institute of Standards and Technology (NIST) in [41], namely, "*an incorrect or suboptimal configuration of an information system or system component that may lead to vulnerabilities*".

3.2.1 Identifying Information Sources. We collected research papers published in the past five years by searching both the Google Scholar and DBLP databases with the *Kubernetes* and *security* keywords; we specifically ensured full consideration of works published in top-tier conferences on both security (e.g., ACM CCS, IEEE S&P, USENIX Security, and NDSS) and networking (e.g., CoNEXT, SIGCOMM, NSDI, and INFOCOM). Regarding industry sources, we went through security standards, whitepapers, and the Kubernetes documentation. In particular, we first considered all industry artifacts cited on the previously identified academic works: the CIS benchmark [33], NSA CISA "Kubernetes Hardening Guidance" [61], OWASP [62], the Microsoft Threat Matrix for Kubernetes [53], the CNCF 2022 Annual Survey [28]. We then extended this corpus with public artifacts created in the last five years by Kubernetes-focused companies: VMWare "State of Kubernetes" [78], Red Hat "2024 State of Kubernetes Security Report" [67] and "Cryptojacking

Table 1. Identified network misconfigurations in Kubernetes, along with their security issues and possible attacks. Section 3.5 discusses their mitigation.

ID	Description	Issue	Possible attack(s)
M1	Port open on container is not declared	Listening on all interfaces by default	Command and control Sensitive port information
M2	Container allocates dynamic ports	Dynamic ports cannot be controlled	Loosened security policies
M3	Port declared on container is not open	Missing checks on declared ports	Data interception / spoofing Data exfiltration
M4A	Compute unit collision	Missing checks on label collision	Man in the middle Server impersonation
M4B	Service label collision		
M4C	Compute unit subset collision		
M4*	Global label collision		
M5A	Service targets unopened port	Missing checks on declared ports Missing checks on existence of target label	Data interception Data spoofing Denial of service Bypassing security checks
M5B	Service targets undeclared port		
M5C	Headless service port is not available		
M5D	Service without target		
M6	Lack of network policies	No isolation between containers	Data interception / spoofing Privilege escalation
M7	Container binds to host network	Network policies do not apply to host	Bypassing network controls

Attacks in Kubernetes" [30], and Aquasec "The State of Kubernetes and Docker Security in 2021" [5]. We finally identified the software tools referenced in both academic and industry sources [13, 67].

3.2.2 Inclusion and Exclusion Criteria. We limited our attention to academic and industry artifacts specifically addressing network-related security issues or lateral movement in Kubernetes, and disregarded those that do not treat the topic in adequate detail or focus on other orchestrators, including related services offered by cloud providers [5, 18, 28, 33, 67, 78]. As for software tools, we restricted our focus to those with each of the following properties: they perform static or runtime analysis or are continuous monitoring tools (e.g., security platforms); they either support or specifically target Kubernetes; they are able to identify at least one type of issue related to networking. We specifically excluded: CNI plugins or tools that depend on a specific cloud provider, such as Tetragon [37] and Azure Netpol manager [54]; policy engines or tools for compliance, such as Kyverno [48] and Open Policy Agent [26]; software tools that are not publicly available, including AutoArmor [50] and Bastion [59]; solutions specifically targeted for Infrastructure as Code, such as Tfsec [9] and Tflint [76]. As a result of this process, we obtained a list of research / industry artifacts [10, 13, 15, 51–53, 56, 61, 62, 70, 80] and software tools [3, 4, 8, 17, 43, 47, 65, 71, 73, 75, 80].

3.2.3 Labeling Misconfigurations. The resulting set of sources was leveraged to create a curated list of security issues broadly related to networking according to the principles of inductive thematic analysis [14]. For this purpose, two of the authors independently labeled security issues into specific codes (namely, instances or types) of misconfigurations. Specifically, they started from an initial list of codes and compared them against the Kubernetes documentation, so as to identify gaps and opportunities to discover novel network misconfigurations. The latter entailed the creation of software tools as well as proof of concept attacks to selected Kubernetes applications (see Section 2.1 for more details). The outcome of this process was an initial set of labels for different misconfiguration types, which were then refined and consolidated in the list presented next.

3.3 Identified Misconfigurations

This section details the identified misconfigurations (each associated with an identifier for convenience), including their possible impact on cluster security (Table 1). Those already recognized by the existing literature are referred to as *Service without target* (*M5D*), *Lack of network policies* (*M6*) and *Container binds to host network* (*M7*); the rest were discovered through independent research.

M1: Open ports are not declared. It occurs when a container actually has ports open (at runtime) but these are not declared in the resource configuration (the Kubernetes YAML files). Open ports are available to all the pods by default, even when no access the host network is allowed, unless restrictive network policies are in place. As a consequence, possible attacks include access to sensitive information as well as command and control. Representative applications exhibiting such security issues are those exposing administrative / debug ports and distributed services leveraging the worker-controller architectural pattern.

M2: Dynamic ports. Dynamic (also referred to as ephemeral) ports are communication endpoints whose identifier (i.e., number) is not explicitly defined by the developer. Instead, the operating system takes care of assigning a port number, which is chosen differently each time the associated server is started. Even though dynamic ports are not a misconfiguration per se, they cannot be specified as such in the declarative configuration of a Kubernetes application. For this reason, we treat them as a network misconfiguration. Even worse, dynamic ports cannot be completely blocked by enforcing network policies unless wide port ranges are specified. Sample applications using dynamic ports are those utilizing control / signaling channels (e.g., by multimedia protocols) or reverse TCP connections.

M3: Declared ports are not open. It occurs when a certain port is specified in the declarative configuration of a container but the application is not actually listening to that port. Container ports can be arbitrarily opened and no check is performed on them by the orchestrator or at the lower layers. Combined with other weaknesses, this misconfiguration allows an attacker to carry out data spoofing, data interception, and data exfiltration just by listening to the declared port. Representative cases are applications supporting different deployment modes (for instance, running in a cluster or stand-alone) as well as scalable services that run in a single instance.

M4: Label collisions. They happen when unrelated resources are identified or targeted by the same set of labels. Collisions can be further classified based on the nature of the objects the labels are applied to. A **Compute unit collision (M4A)** takes place when the same label is applied to different compute units. Moreover, a **Service label collision (M4B)** occurs when multiple services target the same compute unit. Conversely, a **Compute unit subset collision (M4C)** involves unrelated compute units sharing some common labels which are selected as targets for a single service. Collisions also affect different applications deployed into the same cluster; we refer to this case as a **Global label collision (M4*)**. Unfortunately, Kubernetes has no built-in functionality to prevent these: incorrect labeling possibly results in bypassing network policies and causes man-in-the-middle or server impersonation attacks. Sample applications exhibiting label collisions are those that use generic labels to describe resources.

M5: Service with incorrect references. It occurs when a service mistakenly refers to a given compute unit port; it does not generally pose security risks alone, however, it can amplify the reach of an attacker in the presence of other misconfigurations. There are different types of incorrect references, leading to a more detailed classification. One case is when a **Service targets an unopened port (M5A)**; it amplifies the spoofing or interception attack involving *M3* because services are the preferred way to contact applications (rather than compute units). This misconfiguration can

also cause a Denial of Service (DoS) as each request routed to that service will not receive a reply. Another type of misconfiguration manifests when a **Service targets an undeclared port (M5B)**. This might represent an evasion technique to circumvent static security checkers, for instance, to hide an SSH port. Moreover, there is the chance that a **Headless service port is not available (M5C)**. A headless service is simply a DNS record that points to specific compute units in the cluster: a misconfiguration takes place when the service port differs from the one that opens in a container, resulting in mounting DoS (similar to *M5A*). Finally, **Services without target (M5D)** occurs when a service does not match with any compute unit. This implies that every request involving such a service results in a failure. Similar to *M4**, a malicious user with access to the cluster could impersonate the service by deploying a compute unit with matching labels.

M6: Lack of network policies. We consider missing network policies as a misconfiguration, in alignment with the existing literature [33, 72]. Helm charts describing application deployments can specify network policies, but they can be enabled or not. In this case, we consider network policies that are available but not enabled as a misconfiguration.

M7: Container binds to the host network interface. It occurs when the container bypasses the isolation provided by the Kubernetes networking layer and exposes ports to the underlying host network. This happens when the `hostNetwork` field in the declarative specification of a compute unit is set to `true`. As a consequence, the network namespace of the compute unit network becomes bound to the host network namespace, bypassing any network policy attached to it.

3.4 Causes and Consequences

The root cause of *M1* and *M3* is that the declarative nature of the ports is purely documentative, thus, it is not enforced by Kubernetes. This behavior aligns with the Kubernetes networking model; however, inaccurate characterization of ports is a major source of confusion for cluster administrators, especially for those reusing third-party charts. In fact, they would expect the configuration to be authoritative, to the point that they may use it to automatically setup network policies. Unfortunately, these would not properly work in presence of the two misconfigurations. Even worse, ports that are not accurately documented may affect the creation of network policies. Similar considerations apply to *M5A* and *M5B*. Somewhat related, the key issue behind *M2* is that dynamic ports are incompatible with the design principles behind Kubernetes networking [46]. In fact, these ports cannot be practically specified in the configuration of a Kubernetes application.

The cause of *M4* and all its variants is that Kubernetes lacks a mechanism to detect or prevent duplicate labels / selectors across compute units and services. Similarly, *M5C* and *M5D* occur because the orchestrator does not provide any warning that there is no target or port available for a given service. The main reason behind *M6*, instead, is the default "allow all" connectivity policy in Kubernetes, which is clearly too permissive. In contrast, most Linux distributions and major cloud providers generally employ a "deny all" approach. Clearly, attackers can easily move laterally in the cluster if network access is not restricted.

Finally, the main motivation for *M7* is the need to access hardware or operating system-related resources. This typically occurs when exposing metrics, as the Node Exporter component of Prometheus [63] does. Another example is given by scenarios that demand high efficiency, for instance, bandwidth-intensive applications such as those involving machine learning tasks [42]. Using the host network results in making network policies defined for a pod ineffective.

3.5 Mitigation

The last phase of our study entails deriving practical guidelines to avoid misconfigurations and therefore reduce their impact on possible security issues. We now discuss them in detail.

The mitigation of misconfigurations *M1*, *M3*, *M5A* and *M5B* is similar: it involves editing the configuration files to declare all the ports that are open in a container and to ensure that only such ports are bound to services. In particular, Helm charts demand special attention to handle ports that are open depending on specific parameter values (e.g., when optional features are enabled in the applications). As for *M5C*, it is enough to remove the port settings since headless services do not support such a feature. Handling misconfiguration *M2* is problematic, as it requires knowledge of the underlying application. The main objective is to declare the dynamic ports in some way. This can often be achieved by setting specific parameters in the applications to assign static ports instead of using the default (random) values. If this is not possible, the developer of the configuration should add a comment notifying their users that a certain application uses dynamic ports. This is needed because tools relying on network traffic analysis may erroneously generate policies for such ports, as they change every time the pod restarts. All *M4* variants can be addressed by making the labels unique. However, the process requires understanding the relationship between the different application components and the actual purpose behind using the same labels for them. Instead, *M5D* can be fixed by ensuring that each service has a selector matching the labels[1] of a compute unit. Mitigating *M6* requires defining / enabling network policies in the chart, making sure that each policy selects at least one pod and that rules only allow necessary connections. Finally, *M7* can be only remediated by setting `hostNetwork` to `false` after checking that such a change does not result in loss or functionality or performance. In the latter cases, alternative solutions should be sought or at least an in-depth security audit of the relevant pods should be carried out.

4 Evaluation

This section evaluates the proposed approach along different dimensions. First, it considers diverse datasets and carries out an analysis of the findings from running our solution. Second, it compares the effectiveness of the proposed approach against software tools in the state of the art.

4.1 Sources

We considered 287 open source Kubernetes applications defined as Helm charts by six different organizations. Accordingly, we divided them into different datasets as detailed next.

4.1.1 Datasets. We selected publicly available Kubernetes applications in the form of Helm Charts. We only picked applications from reputable sources and actively maintained. We also chose applications within specific organizations to represent diverse use cases, to gain a broad view of misconfigurations related to lateral movement. We classify and describe the datasets next.

- *Sharing*. This category is given by charts defined by well-known organizations such as Bitnami (part of VMWare) and Banzai Cloud (owned by Cisco) which develop and officially support configurations meant to be reused and shared.
- *Internal*. This class is represented by organizations maintaining Helm charts for their own software. Specifically, we considered Wikimedia and the European Environmental Agency (EEA) as representative examples.
- *Production*. The last group is represented by organizations such as the Cloud Native Computing Foundation (CNCF) and Prometheus Community which develop configurations for their own applications that are purposely built for production use. This use case is also representative of applications that are jointly deployed into the same Kubernetes cluster.

4.1.2 Considered applications. We selected only applications that run in their default configuration or with minimal changes for each dataset. Therefore, we did not consider applications that require

[1]This can be checked, for instance, with the `kubectl get pods -l <svc_selector>` command.

Table 2. Breakdown of network misconfigurations by dataset. In total, 634 misconfigurations were found.

					M4				M5					
Dataset	**Affected apps**	**M1**	**M2**	**M3**	**M4A**	**M4B**	**M4C**	**M4***	**M5A**	**M5B**	**M5C**	**M5D**	**M6**	**M7**
Banzai Cloud	51 / 51	13	2	17	8	4	0	0	0	2	0	0	51	0
Bitnami	158 / 158	106	26	40	25	10	0	5	2	14	3	0	156	7
CNCF	7 / 10	10	0	4	0	0	0	0	6	0	0	0	7	0
EEA	8 / 19	7	0	1	0	1	0	0	0	0	0	0	0	0
Prometheus C.	25 / 25	42	4	3	0	0	0	0	1	4	0	0	25	4
Wikimedia	10 / 27	10	3	2	2	1	1	0	2	1	0	0	2	0
Total	259 / 287	188	35	67	35	16	1	5	11	21	3	0	241	11

specific environment variables, cloud providers, or images from private container registries. This is necessary to automate the testing process and to ensure the accuracy of the runtime analysis.

4.2 Setup

We now briefly discuss the implementation of a solution that is able to identify the misconfigurations presented in the previous section. We also describe how we handled special cases of importance.

4.2.1 Implementation. We employed the characterization of the misconfigurations in Section 3.3 to create a list of machine-readable rules. We take a Helm chart as input then carry out both static and runtime analysis. We then combine the obtained results and evaluate them against the machine-readable rules. In doing so, we consider both the labels and the selectors associated with pods and services in addition to discrepancies between declaration and runtime behavior. Specifically, we carry out static analysis through a custom software that parses the YAML files and extracts the relevant information, including container ports, service ports, labels, and selectors. For the runtime analysis, we install each application into an empty Kubernetes cluster and observe its runtime behavior by following the approach in [15]. In detail, we used Minikube [25] version 1.23 and Kubernetes version 1.25. We opted for a virtualized environment to automate the process and ensure isolation. In particular, we delete and recreate the cluster after analyzing individual applications to guarantee that unrelated resources do not affect each other. Once all applications have been individually evaluated, we search for cluster-wide misconfigurations by checking the labels and selectors of every single application.

4.2.2 Special Cases. We also took additional steps to address a few corner cases that require special handling. The first one is related to the dynamic ports associated with *M2*. These ports change every time the application is started, therefore, they are not captured by a single snapshot provided by the runtime analysis. To address such an issue, we perform the analysis twice then compare the ports detected in the two distinct iterations. We report *M2* when the detected ports differ. The second one occurs for *M7* because the application has access to the host network. In such a case, the runtime analysis reports all the ports that are open at the host, possibly including those of other components unrelated to the application. We address this issue by carrying out a preliminary analysis of the open ports in the host, which are then removed from the final output.

4.3 Analysis

The misconfigurations found by analyzing the considered applications are detailed next.

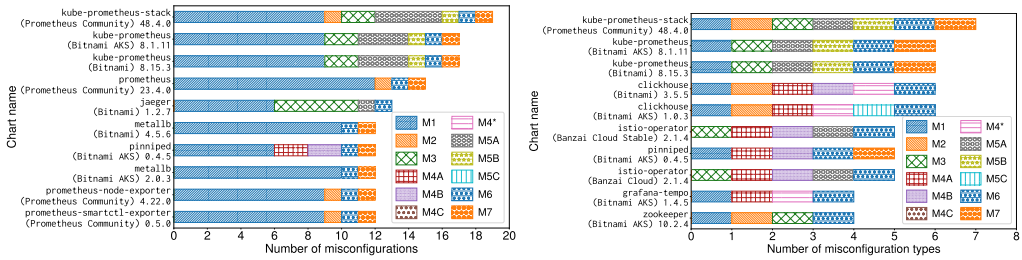

Fig. 3. The ten applications with the highest number of (a) misconfigurations and (b) misconfiguration types.

4.3.1 Results. Table 2 details the outcome of our analysis, divided by dataset. As a result, 90% (259 out of a total of 287) of the applications in all the datasets have one or more misconfigurations. The most common misconfiguration types are *Lack of network policies* (*M6*), *Port not declared* (*M1*) and *Port not open* (*M3*). These misconfigurations originate from lacking or inaccurate network boundaries on the Kubernetes resources, which could harm the cluster integrity. Approximately 10% of the considered applications listen to ephemeral ports *(M2)* and a similar number have label collision issues (*M4*). Among the label collision issues, the compute unit collision (*M4A*) is the most common, with 35 instances discovered. Service misconfigurations (*M5*) occur but less frequently; they only affect individual applications, not the cluster (those indicated as *M4** in the table). Finally, only 3% of applications have containers binding to the host network namespace (*M7*). Figure 3a shows the ten most misconfigured applications and Figure 3b shows the ten applications with the highest number of different concurrent misconfigurations. These are all distributed applications spanning from log collection to load-balancing. Among the most misconfigured applications by number, all lack network policies (*M6*) and have multiple open ports that are not declared (*M1*); 9 out of 10 bind to the host network namespace (*M7*); 4 out of 10 applications listen to ephemeral ports (*M2*). Regarding the applications with the most types of misconfigurations, they also lack network policies but the other misconfigurations are more evenly distributed.

We then focused our attention on misconfigurations based on the type of dataset they belong to. For this purpose, we grouped the datasets according to their use case as described in Section 4.1. We noted that every application in the *sharing* and *production* datasets exhibited at least one misconfiguration – their average number of misconfigurations per application was 3.35 and 4.44, respectively. In contrast, misconfigurations were found in only 39% of charts built for *internal* use, with an average number of misconfigurations slightly above one (specifically, 1.11). These results show that charts built by third parties have a significantly higher number of misconfigurations than those built for internal use: three to four times on average. This happens as charts are shared to promote adoption; as a consequence, their creators favor support for diverse functionality (e.g., optional application components) over strict enforcement of secure defaults. In contrast, charts intended for internal use are tightly integrated, whereas those for production environments generally rely on additional components or tooling. See also Section 5.2 for additional discussion.

We lastly characterized how misconfigurations are distributed across the different applications in the datasets (Figure 4a). The data from our experiments revealed that 5% of the applications exhibit at least 10 misconfigurations each, overall accounting for 25% of the total number of misconfigurations. Moreover, 8% of the applications had between 5 and 9 misconfigurations each, corresponding to about 22% of the total. The rest of the misconfigurations were almost evenly distributed between the remaining applications. We have found that monitoring applications (such

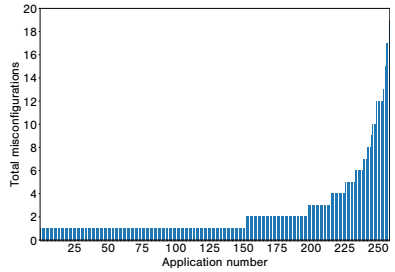

(a)

Dataset	Network Policies Enabled	Network Policies Affected	Reachable Pods	Reachable Services
Bitnami	48 (2)	3	14 (1)	0
CNCF	4 (0)	0	0 (0)	0
EEA	19 (19)	8	13 (3)	1
Prometheus C.	5 (0)	3	32 (3)	0
Wikimedia	25 (23)	4	8 (5)	2

(b)

Fig. 4. (a) Total misconfigurations per application. (b) Impact of network policies on endpoint reachability.

as observability frameworks) often exhibit open ports that are not declared (*M1*). Furthermore, applications that leverage a master-worker architecture (for instance, workflow engines) generally rely on dynamic ports (*M2*) for coordination purposes. Finally, applications providing a substantial number of microservices (including those for federated services) tend to have declared ports that are not opened (*M3*), because most of the corresponding features are not actually enabled at runtime.

4.3.2 Impact of Network Policies. Network policies can be ill-defined, leading to a false sense of security. To evaluate their impact on misconfigurations, we analyzed all charts that define network policies and enabled them if they were not active by default. We followed the same methodology described in Section 4.2 to determine whether network misconfigurations were mitigated by the policies. Specifically, we parsed the results of the runtime analysis and searched for endpoints corresponding to misconfigured ports that remained reachable from within the cluster. Figure 4b shows the obtained results; the Banzai Cloud dataset is not reported as none of the applications therein had any network policies. Misconfigured charts with loose network policies are indicated as affected in the table. The values within parentheses indicate the use of dynamic ports.

The table shows that enabling network policies did not remedy misconfigurations in most cases. The allowed connections to misconfigured pods are higher than those to misconfigured services, which is aligned with the previous findings and also due to the presence of dynamic ports. Clearly, the sheer amount of reachable pods and services depends on the size of the applications, in terms of the number of containers therein. Also note that the results are obtained by deploying a single application in the cluster; the opportunities for reaching misconfigured ports would increase for multiple applications deployed at once. In summary, the datasets contain unwanted or unknown connections even when network policies are applied. This is not always due to dynamic ports, but it is also caused by erroneous settings. For instance, we noticed cases where strict policies targeted pods with access to the host network; however, network policies are not effective in such a case.

4.4 Comparison with State of the Art

This section evaluates our solution against state-of-the-art security tools for Kubernetes.

4.4.1 Considered Solutions. We considered representative tools among those found as described in Section 3.2. For clarity, we divide them in the categories detailed below.

- *Static analysis.* They analyze Kubernetes manifests, YAML files, and other static configurations before they are deployed in the cluster, comparing them against a database of best practices [17, 47, 65] or security guidelines [71, 80], such as the CIS benchmark for Kubernetes [33]. They also offer actionable insights to improve the security of the deployments.

Table 3. Network misconfigurations detected by the considered tools and our solution. The symbols indicate whether these were ● found, ◐ partially found (i.e., either find less misconfigurations or require multiple runs), × missed, or — not applicable (i.e., could not be found by intrinsic limitations in the type of tool).

						M4				M5					
Tool	**Version**	**Type**	**M1**	**M2**	**M3**	**M4A**	**M4B**	**M4C**	**M4***	**M5A**	**M5B**	**M5C**	**M5D**	**M6**	**M7**
Checkov [17]	3.2.23	Static	—	—	—	×	×	×	—	—	×	×	×	●	●
Kubeaudit [71]	0.22.1	Static	—	—	—	×	×	×	—	—	×	×	×	●	●
KubeLinter [43]	0.6.8	Static	—	—	—	×	×	×	—	—	×	×	●	×	●
Kube-score [80]	1.18.0	Static	—	—	—	×	×	×	—	—	×	×	●	●	×
Kubesec [47]	2.14.0	Static	—	—	—	×	×	×	—	—	×	×	×	×	●
SLI-KUBE [65]	N/A	Static	—	—	—	×	×	×	—	—	×	×	×	×	●
Kube-bench [3]	0.7.1	Runtime	×	×	×	×	×	×	—	×	×	×	×	×	●
Kubescape [8]	3.0.3	Hybrid	×	×	×	◐	◐	◐	×	×	×	×	×	●	●
Trivy [4]	0.49.1	Hybrid	×	×	×	×	×	×	×	×	×	×	×	×	●
Neuvector [75]	5.3.0	Platform	×	×	×	×	×	×	×	×	×	×	×	×	●
StackRox [73]	3.74.9	Platform	×	×	×	×	×	×	×	×	×	×	×	×	●
Our solution	—	Hybrid	●	●	◐	●	●	●	●	●	●	●	●	●	●

- *Runtime analysis.* Analyze Kubernetes resources deployed in a cluster via the Kubernetes API [3] or by inspecting running applications [15], providing a comprehensive view of the system based on all available cluster resources. Some tools [7] impersonate malicious actors to assess application security.
- *Other approaches* A few tools support both static and runtime analysis [4, 8], making them *hybrid* solutions. Another class of tools is represented by those performing continuous security monitoring, which are also commonly known as *security platforms* [73, 75].

4.4.2 Methodology. We selected representative Kubernetes configurations exhibiting the misconfigurations in Table 1 for the evaluation by ensuring that the considered charts are representative of all the misconfiguration types identified in the paper. We employed this approach for convenience, to reduce the input provided to the different tools and to simplify the analysis of the issues reported in the output. We applied the configuration in a running Kubernetes cluster with the same setup as Section 4.1 for the tools belonging to the runtime, hybrid, and platform categories.

Note that specific tools are unable to identify certain misconfigurations as a limitation intrinsic to their nature. For instance, tools only performing static analysis cannot discover issues occurring at runtime; conversely, those purely carrying out analysis at runtime are unable to detect misconfiguration due to cluster-wide collisions. We indicate these as "not applicable" in the table.

4.4.3 Results. Table 3 shows the misconfigurations identified by the considered tools, grouped by their type for convenience. First of all, we can see that none of the existing tools are able to at least partially recognize all misconfigurations, even though they successfully detect one or more of them – the lack of network policies (*M6*) and host network mapping (*M7*) are the most recognized.

Clearly, tools only supporting static analysis are (by their inherent nature) unable to identify most of the misconfigurations because they only inspect the configuration files. Even so, they still miss many of the misconfigurations they could possibly detect: all of those involving label collision (*M4*) and most of what relates to services (*M5*) – only KubeLinter [43] and Kube-score [80] identify services without target (*M5D*). In contrast, security platforms and hybrid tools could in principle recognize most of the misconfigurations considered in this work, as they also analyze the

applications at runtime. However, their effectiveness is poor in practice: only Kubescape [8] reports deployments in which common labels are used for resources, hinting at issues involving label collisions. Security platforms do get system information from the Kubernetes APIs and monitor processes and network sockets; however, they do not make any effort in notifying the user about potentially misconfigured resources, effectively delaying the detection of lateral movement after it has already occurred. In contrast, our approach was the only one to identify all the misconfigurations.

We also checked the source code of the considered tools to obtain more insights why misconfigurations were missed. We found that most of them do not check the relationship between Kubernetes resources of different types, for instance, between a service and a pod. This explains why static analysis tools were not able to identify label collisions and service-related issues as we instead do. Moreover, tools performing runtime analysis only query the configuration of resources from the API Server but do not actually inspect the runtime environment of the containers (for instance, open ports). Finally, most security platforms do not check for network-related settings but allow to record the entire traffic and possibly generate network policies based on it. Unfortunately, the latter is not enough as it can only be done assuming that all recorded traffic is intended and legitimate.

5 Disclosure

Our approach allows to identify diverse misconfigurations, however, it does not assess their severity. This is because such an assessment requires expert knowledge in the domain of the application as deployed in the Kubernetes cluster. For this reason, we reported the results of our analysis to the relevant organizations and engaged with their developers to understand the actual security implications of the misconfigurations in specific cases.

5.1 Process and Follow-up

We first performed a responsible disclosure of our findings to the organizations associated with the datasets we analyzed (see Table 2), reporting a total of 634 misconfigurations. We then engaged in an active and constructive discussion with these organizations so as to assess their severity through expert knowledge in the domain of the application. Appendix A.1 provides additional details on the methodology used for the disclosure.

5.1.1 Acceptance and Severity Assessment. All developers recognized the issues we reported as misconfigurations. Those using software tools for security analysis also acknowledged that the issues identified in our work had not been detected before. Moreover, the respondents reported label collisions (*M4*) as the most critical and "declared port not open" (*M3*) as the least critical in terms of security risks. A few developers revealed that security issues related to lateral movement were internally classified as minor. This substantiates our premise that cluster-internal networking is a security threat that is often overlooked.

Regarding the severity of the reported misconfigurations, a respondent from the European Environment Agency explicitly stated in their feedback that the changes were security mitigations, whereas a developer from Bitnami mentioned "*improved security of the application*" as a benefit in their pull requests. Finally, a respondent from Wikimedia reported the outcomes of a security review of the reported misconfigurations; their feedback excluded the chance of remote code execution, but did not rule out other attacks. Overall, two of the respondents addressed (at least partially) the identified misconfigurations in less than 72 hours after receiving the report – at the time of writing, misconfigurations affecting 30 applications have been fixed across five different organizations. Appendix A.2 provides more details on the specific changes that fixed these misconfigurations.

5.1.2 Accuracy and Analysis of Special Cases. Engaging with the developers also allowed us to detect false positives in our implementation, as well as to identify distinctive configurations for

certain applications. One of the developers pointed out issues on UDP ports displayed as part of the misconfigurations. After investigating these, we found out that random UDP ports were sometimes reported as part of the runtime analysis. This kind of false positives amount to about 8% of the total misconfigurations initially identified; they are not included in the final results summarized in Table 2. We also received feedback regarding the service label collision misconfiguration, explaining a case where the collision was deliberate. It involved two services targeting a pool of database pods, wherein a service load-balances the requests to the pool of pods while the other one (headless) always retrieves the same instance, intended for write operations.

A special use case was found in the Wikimedia dataset. Our analysis of the dataset showed that nearly half of the applications expose ports that are not declared. A careful inspection of the repository and of the tooling used by Wikimedia revealed a rather unconventional use of the port declaration, which we confirmed with the developers. In fact, Wikimedia charts only declare the ports of the applications that are needed for communication with other components, whereas undeclared ports are only intended to be used within the component. This is enforced with a custom tool that automatically generates network policies from the ports declared in the configuration files. The misconfigurations we identified in this scenario cannot actually be considered false positives but show how users rely on the declarative configuration of applications, as discussed in Section 3.4.

5.2 Feedback

This section summarizes the findings of our analysis that emerged as part of the disclosure process.

5.2.1 Implicit Security Assumptions. The recipients of the disclosures reacted differently to the reported misconfigurations. Organizations that create charts for third parties stated that the port information included in the chart description should not be used to create network policies. They rely on the assumption that users leveraging their charts create their network policies by manually inspecting connections with real traffic, as one of the respondents stated: "*Security sensitive users should probably start with 'allow nothing' and add allowances based on real world traffic policies*". The respondents agreed that network policies are needed. Those who did not include network policies with their applications justified their choice because of their peculiar features: "*I agree users probably should write network policies for their applications [...] I am not very convinced they should be in the helm chart, as they are so specific to each user*". They also suggested using external tools (such as third-party network plugins) to define and enforce network policies.

It is also interesting to notice how organizations that consume charts created by others value accurate declaration of port information, compared to those that create charts for their own applications. This is because they use such information to generate network policies. We are convinced that omitting network policies as part of applications because of the above-mentioned concerns is not only unsafe but fundamentally incorrect. Kubernetes, indeed, treats most of the network resources as interfaces, so that users can choose how to implement them. Therefore, the `NetworkPolicy` resource is a suitable candidate to provide a generic policy description that can later be implemented by other plugins.

5.2.2 Information Hiding. We have also demonstrated that users actually leverage the ambiguous nature of the port description to fit their own needs. This happened for Wikimedia, where the port description in the pods is used to automatically generate network policies. This case is interesting because it follows the principle of "information hiding", where an interface is created between modules, exposing public interfaces and hiding those that are private.

6 Discussion

This section summarizes our findings, offers recommendations on networking-related best practices for cloud-native applications running in Kubernetes, and examines the limitations of our work.

6.1 Findings

Ambiguous or Unspecific Guidelines. Following security best practices may create a false sense of security for both organizations and individual developers. This is especially because these best practices are not evenly developed. For instance, the CIS Kubernetes benchmark [33] contains 30 recommendations on how to secure the API server and 13 about role-based access control, but only two of them refer to network policies. To make it worse, these guidelines are ambiguous or not specific. In fact, they only state "ensure that the [container networking interface] supports Network Policies" (5.3.1) and "ensure that all Namespaces have network policies" (5.3.2). Although the recommendations are indeed valid, they fail to provide enough detail on how to correctly configure network policies. Moreover, many key aspects of internal networking are left out, including how to ensure that all the resources are correctly isolated, and it is the duty of the administrator to manually check them. As previously highlighted, Kubernetes does not automatically alert if a resource to be deployed has a label collision with an existing one in the cluster, or if the ports specified in a network policy actually match those of a given pod or service.

Reliance on Security Tools. Organizations rely on security tools to assess the security of their Kubernetes clusters [67, 78] and to alleviate the manual work required to secure the cluster. However, these tools often automate the checks suggested by the guidelines, inheriting their limitations and ambiguities. Therefore, they are able to correctly find inconsistencies on access control policies but miss most of the networking misconfigurations presented in this paper. Unfortunately, security tools are still immature [56] and do not completely cover all the spectrum of misconfigurations that can appear in a complex environment such as a Kubernetes cluster. Therefore, we recommend developers not to fully depend on those tools but to actively incorporate the measures proposed here to assess the security of their applications and clusters.

Zero Trust and Service Meshes. Zero trust is an emerging approach to security that moves away from the traditional concept of static network perimeter [69]. In fact, it assumes that no implicit trust is granted to resources purely based on the physical / network location and ownership of a device. As a result, all communication flows should be strongly authenticated and authorized [66]. In practice, zero trust in Kubernetes is achieved by employing a service mesh such as Istio [23] – a software layer that allows to monitor and control the information flow between different microservices, typically through sidecar proxies that authenticate, authorize and encrypt communications [49]. The misconfigurations presented in this work can bypass the security mechanisms of the service-mesh. In fact, using a service mesh does not imply that the microservices in the application can *only* be accessed through the mesh, since Kubernetes allows unbounded pod-to-pod communication by default. Moreover, if applications in the mesh are misconfigured, the related security issues will inevitably concern the mesh itself as well. Therefore, network policies are still needed to protect direct access between components and enable defense in depth [38].

6.2 Recommendations

End users. Users need to understand the network properties and settings of the applications they deploy to ensure proper network segregation. This involves carefully reviewing the network ports associated with services and pods, as well as their relationships through labels and selectors. In addition, users should verify that network policies are available, enabled, and properly configured.

A recommended practice is to use runtime analysis tools [15, 75] to monitor open ports in the cluster and critically assess their necessity – especially for optional or non-essential components.

A special case involves the use of dynamic ports. Network policies cannot handle them, therefore, users should block access between a pod with dynamic ports and any unrelated service. A better alternative is to change the configuration to prevent the use of dynamic ports. It is often sufficient to manually configure the ports by setting environment variables or editing the default configuration.

Chart creators. Creators should include all necessary network-related information in their charts for users to effectively deploy them. As a first step, they should document all ports that are open in the pods and services of their applications. Moreover, it is advisable not only to define but also to enable network policies by default in the chart, thereby establishing a more secure baseline for Kubernetes applications. The best way to add network policies is to use the default `NetworkPolicy` resource because it does not rely on third-party plugins. Finally, creators should carefully define templates and default parameters in their charts to ensure secure and reliable deployments.

6.3 Limitations

Several factors may affect the accuracy and validity of our results. Some of them are due to our software implementation; in particular, the runtime analysis erroneously reports UDP ports as open in some cases (see Section 5.1.2). Moreover, it might miss ports that are not open during the analysis, for instance, those triggered by incoming traffic or due to port knocking techniques.

Our approach assumes that all network-related configuration is self-contained in a Helm chart, particularly, that it is defined through core Kubernetes components. However, users may rely on external tools or third-party Kubernetes resources (e.g., network plugins), as highlighted in Section 5.1.2. Another limitation of our approach is the lack of a ground truth. As a result, it cannot provide a conclusive security assessment of an application without expert knowledge (see Section 5.1.1). Instead, our analysis reports misconfigurations similar to what linters do.

Finally, our evaluation methodology relies on automatically deploying applications defined as Helm charts. However, some applications require manual configuration or rely on cloud provider-specific components to be successfully deployed; we could not consider these in our evaluation. Despite covering several diverse datasets, the results we obtained may not generalize to other scenarios, for instance, those not involving open-source applications.

7 Related work

Defensive security in the cloud spans multiple layers of the network infrastructure [32, 39]. The rest of this section focus on the existing literature that is more closely related to our work.

Network isolation for containers. Nam et al. [59] carry out a security analysis of container networking and devise a high-performance network stack that enforces security through a communication sandbox. Similarly, Nakata et al. [58] propose a sandboxing mechanism to improve network isolation for containers with a low overhead. In contrast, this work allows to harden network security of Kubernetes clusters without relying on additional components. Zhu and Gehrmann [81] design and implement a system to generate AppArmor profiles for Kubernetes applications, including network rules, by collecting measurements at runtime. Their system requires legitimate traffic to produce meaningful results, whereas our work allows to identify unnecessary or malicious connections in a cluster purely based on applications' endpoints. Open Policy Agent [26] is a general-purpose engine for policy enforcement; it offers native support for Kubernetes as an admission controller that intercepts requests to the orchestrator API and applies policies. Service meshes [49] such as Istio allow to easily realize access control, encryption, and end-to-end

authentication in a Kubernetes cluster. However, these solutions rely on correct policies (as noted in Section 6.1), whereas our solution enables finding misconfigurations that can be easily overlooked.

Attacks to Kubernetes clusters. Minna et al. [55] point out that the network abstractions employed by Kubernetes may result in unexpected attacks when users approach cloud security with a mental model derived from physical networks. As a step further, this work defines different types of misconfigurations and identifies them in real applications belonging to different use cases. Yang et al. [79] demonstrate lateral movement attacks caused by excessive permissions on Kubernetes service accounts. Gu et al. [34] devise a solution to automatically detect such permissions that can be exploited to harm a Kubernetes cluster. In contrast, we rather focus on network misconfigurations that originate from Kubernetes abstractions rather than on loosely defined role-based access control rules. Ben David and Bremler-Barr [12] show that Kubernetes clusters are prone to Economic Denial of Sustainability (EDoS) attacks – those causing economic damage through unnecessary use of resources – triggered by its auto-scaling mechanism. Chamberlain et al. [16] devise a model based on Markov decision processes to characterize EDoS attacks to Kubernetes. Our work does not address such a specific attack but considers network misconfigurations that can lead to diverse security issues, depending on the specific behavior of a Kubernetes application and its deployment.

Security misconfigurations. Several studies address system configuration in the context of infrastructure as code [64]. Similarly, Dietrich et al. [31] analyze the causes behind the occurrence of misconfigurations by surveying system operators. In contrast, our work specifically targets Kubernetes and does not rely on specific vendors or cloud providers. Shamim et al. [70] accurately review the best practices for securing Kubernetes clusters. They advocate to change default configurations and set audit / network control policies, however, they do not provide details on how to do that. Instead, our work offers a comprehensive account of network misconfigurations, including guidelines on how to fix them. A report by Red Hat on the state of Kubernetes security based on a large-scale survey [67] has found that 40% of the respondents experienced security issues related to misconfigurations. Unfortunately, the report does not include a detailed analysis on their nature, as we do here instead. Rahman et al. [65] carry out an empirical study of security misconfigurations by manually inspecting a dataset of open-source Kubernetes manifests, then validate their findings through a custom static analysis tool. Instead, we apply both static and runtime analysis, consider multiple datasets, and specifically focus on misconfigurations related to Kubernetes networking.

Analysis of Kubernetes applications. Minna and Massacci [56] survey runtime analysis tools for Kubernetes developed by both the research community and the industry; as a result, they identify orchestration security and resilience to internal threats as open challenges. We specifically target the latter in the context of lateral movement within a Kubernetes cluster. Blaise and Rebecchi [13] devise a graph-theoretic methodology to assess the security of microservice-based applications defined in terms of Helm charts. However, their work derives connectivity purely based on what is declared in the configuration files; instead, our solution also leverages runtime analysis for a more reliable characterization of network misconfigurations. Kubesonde [15] employs runtime analysis to derive the network connectivity of Kubernetes applications. Specifically, it just reports such a connectivity but does not assist users in identifying possible issues related to it as we do in this work. Finally, there is a large amount of software tools that help users find security issues related to Kubernetes [3, 4, 8, 17, 43, 47, 65, 71, 73, 75, 80]. Unfortunately, they do not explicitly point out specific network misconfigurations or fail to identify most of them, as demonstrated by Table 3.

8 Conclusion

This work investigated security issues originated by network misconfigurations affecting Kubernetes clusters, with special focus to lateral movement. A systematic review formed the basis for us to identify novel network misconfigurations that could result in security threats. Accordingly, we analyzed a large amount of open-source Kubernetes applications from several organizations. Our evaluation revealed that application deployments exhibit a significant amount of these misconfigurations. The related findings were disclosed to the relevant organizations, most of which acknowledged and fixed the misconfigurations. Our work presents easily identifiable sources of security issues, thereby helping to secure Kubernetes deployments and prevent practical attacks.

There are still interesting directions for future research. We believe that applications would benefit from being specified through modular charts with clearly defined required / optional dependencies. This would allow to more precisely specify the intended network connectivity between components, which could then be leveraged to possibly derive network policies in a automated way. Indeed, Kubernetes is configuration-centric, yet it is still difficult for end users to understand the actual connectivity between different microservices. We expect visualization and monitoring tools to become more effective in assisting users by explicitly supporting network-related metrics and providing proactive advice on the resulting security issues, for instance, through machine learning.

Acknowledgments

This work was partially supported by: the Research Council of Finland under grants number 345964 and 357533; the INCIBE-UPV's Chair of Cybersecurity funded by the European Union under the NextGenerationEU initiative through the Spanish government's Plan de Recuperación, Transformación y Resiliencia.

References

[1] Yasemin Acar et al. 2017. Developers Need Support, Too: A Survey of Security Advice for Software Developers. In *2017 IEEE Cybersecurity Development (SecDev)*. 22–26. https://doi.org/10.1109/SecDev.2017.17

[2] Amit Seal Ami, Nathan Cooper, Kaushal Kafle, Kevin Moran, Denys Poshyvanyk, and Adwait Nadkarni. 2022. Why crypto-detectors fail: A systematic evaluation of cryptographic misuse detection techniques. In *2022 IEEE Symposium on Security and Privacy (SP)*. 614–631. https://doi.org/10.1109/SP46214.2022.9833582

[3] Aqua 2024. Kube-bench: secure checks to Kubernetes deployments. https://github.com/aquasecurity/kube-bench Accessed on September 25, 2024.

[4] Aqua Security 2024. Trivy: security scanner for Kubernetes. https://github.com/aquasecurity/trivy Accessed on September 25, 2024.

[5] Aquasecurity. 2021. Aquasecurity The State of Kubernetes and Docker Security in 2021. https://www.aquasec.com/news/the-state-of-kubernetes-and-docker-security-in-2021/ Accessed on September 25, 2024.

[6] Aquasecurity. 2023. Aquasecurity Blog. https://www.aquasec.com/news/kubernetes-clusters-under-attack/ Accessed on September 25, 2024.

[7] Aquasecurity. 2024. Kube-hunter: Hunt for security weaknesses in Kubernetes clusters. https://github.com/aquasecurity/kube-hunter Accessed on September 25, 2024.

[8] Aquasecurity. 2024. Kubescape. https://github.com/kubescape/kubescape Accessed on September 25, 2024.

[9] Aquasecurity 2024. TFsec. https://aquasecurity.github.io/tfsec/. Accessed on September 25, 2024.

[10] Kubernetes Authors. 2024. Securing a Cluster. https://kubernetes.io/docs/tasks/administer-cluster/securing-a-cluster/ Accessed on September 25, 2024.

[11] Christian Banse, Immanuel Kunz, Angelika Schneider, and Konrad Weiss. 2021. Cloud Property Graph: Connecting Cloud Security Assessments with Static Code Analysis. In *2021 IEEE 14th International Conference on Cloud Computing (CLOUD)*. 13–19. https://doi.org/10.1109/CLOUD53861.2021.00014

[12] Ronen Ben David. and Anat Bremler-Barr. 2021. Kubernetes Autoscaling: YoYo Attack Vulnerability and Mitigation. In *Proceedings of the 11th International Conference on Cloud Computing and Services Science - CLOSER*. INSTICC, SciTePress, 34–44. https://doi.org/10.5220/0010397900340044

[13] Agathe Blaise and Filippo Rebecchi. 2022. Stay at the Helm: secure Kubernetes deployments via graph generation and attack reconstruction. In *2022 IEEE 15th International Conference on Cloud Computing (CLOUD)*. 59–69. https:

//doi.org/10.1109/CLOUD55607.2022.00022
[14] Virginia Braun and Victoria Clarke. 2021. Thematic analysis: A practical guide. (2021). https://uk.sagepub.com/en-gb/eur/thematic-analysis/book248481
[15] Jacopo Bufalino, Mario Di Francesco, and Tuomas Aura. 2023. Analyzing network-layer access and isolation between microservices with Kubesonde. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '23), December 3–9, 2023, San Francisco, CA, USA.*
[16] Jonathan Chamberlain, Jilin Zheng, Zeying Zhu, Zaoxing Liu, and David Starobinski. 2025. Exploiting Kubernetes Autoscaling for Economic Denial of Sustainability. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 9, 2 (June 2025), 29. https://doi.org/10.1145/3727114
[17] Checkov 2024. Checkov: static code analysis for infrastructure as code. https://checkov.io/ Accessed on September 25, 2024.
[18] Michele Chiari, Michele De Pascalis, and Matteo Pradella. 2022. Static Analysis of Infrastructure as Code: a Survey. In *2022 IEEE 19th International Conference on Software Architecture Companion (ICSA-C).* 218–225. https://doi.org/10.1109/ICSA-C54293.2022.00049
[19] Cloud Native Computing Foundation (CNCF). 2020. *CNCF Helm Project Journey Report.* Cloud Native Computing Foundation. https://www.cncf.io/wp-content/uploads/2020/08/CNCF_Helm_Project_Journey_Report.pdf
[20] CNCF. 2024. ArtifactHub. https://artifacthub.io Accessed on September 25, 2024.
[21] CNCF 2024. CNCF Cloud Native Interactive Landscape. https://landscape.cncf.io/. Accessed: April 16, 2024.
[22] CNCF. 2024. Helm: The package manager for Kubernetes. https://helm.sh/ Accessed on September 25, 2024.
[23] CNCF 2024. Istio. https://istio.io/ Accessed on September 25, 2024.
[24] CNCF 2024. Kubernetes: Production-Grade Container Orchestration. https://kubernetes.io/ Accessed on September 25, 2024.
[25] CNCF. 2024. Minikube. https://minikube.sigs.k8s.io/docs/ Accessed on September 25, 2024.
[26] CNCF. 2024. OPA: Open Policy Agent. https://www.openpolicyagent.org/ Accessed on September 25, 2024.
[27] CNCF 2024. Prometheus: A Next-Generation Monitoring System. https://prometheus.io/ Accessed on September 25, 2024.
[28] CNCFSurvey 2024. CNCF 2022 Annual Survey. https://www.cncf.io/reports/cncf-annual-survey-2022/ Accessed on September 25, 2024.
[29] Concourse 2024. Concourse CI. https://concourse-ci.org/. Accessed: December 2, 2024.
[30] Wei Lien Dang. 2020. Cryptojacking Attacks in Kubernetes: How to Stop Them. Red Hat Blog. https://cloud.redhat.com/blog/cryptojacking-attacks-in-kubernetes-how-to-stop-them Accessed on September 25, 2024.
[31] Constanze Dietrich, Katharina Krombholz, Kevin Borgolte, and Tobias Fiebig. 2018. Investigating System Operators' Perspective on Security Misconfigurations. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security* (Toronto, Canada). 1272–1289. https://doi.org/10.1145/3243734.3243794
[32] Billawa et al. 2022. SoK: Security of Microservice Applications: A Practitioners' Perspective on Challenges and Best Practices. In *Proceedings of the 17th International Conference on Availability, Reliability and Security (ARES '22).* Article 9, 10 pages. https://doi.org/10.1145/3538969.3538986
[33] CIS (Center for Internet Security). 2024. *CIS Kubernetes Benchmark.* Technical Report. Center for Internet Security. https://www.cisecurity.org/benchmark/kubernetes/
[34] Yue Gu, Xin Tan, Yuan Zhang, Siyan Gao, and Min Yang. 2025. EPScan: Automated Detection of Excessive RBAC Permissions in Kubernetes Applications. In *2025 IEEE Symposium on Security and Privacy (SP).* 3199–3217. https://doi.org/10.1109/SP61157.2025.00011
[35] Helm. 2024. Helm Documentation: Charts. https://helm.sh/docs/topics/charts/. Accessed: December 2, 2024.
[36] Aftab Hussain and Anton Burtsev. 2020. *Common Vulnerabilities and Exposures in the Cloud.* Technical Report. University of California at Irvine. https://aftabhussain.github.io/documents/pubs/tech-report20-cve.pdf
[37] Isovalent. 2024. Tetragon. https://tetragon.io/. Accessed on September 25, 2024.
[38] Istio. 2024. Security Best Practices – Defense in depth with NetworkPolicy. https://istio.io/latest/docs/ops/best-practices/security/#defense-in-depth-with-networkpolicy Accessed on September 5, 2024.
[39] Hyunsu Jang, Jaehoon Jeong, Hyoungshick Kim, and Jung-Soo Park. 2015. A Survey on Interfaces to Network Security Functions in Network Virtualization. In *2015 IEEE 29th International Conference on Advanced Information Networking and Applications Workshops.* 160–163. https://doi.org/10.1109/WAINA.2015.103
[40] Omar Jarkas, Ryan Ko, Naipeng Dong, and Redowan Mahmud. 2025. A Container Security Survey: Exploits, Attacks, and Defenses. *ACM Comput. Surv.* 57, 7, Article 170 (Feb. 2025), 36 pages. https://doi.org/10.1145/3715001
[41] Arnold Johnson, Arnold Johnson, Kelley Dempsey, Ron Ross, Sarbari Gupta, and Dennis Bailey. 2011. *Guide for security-focused configuration management of information systems.* NIST Special Publication 800-128. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-128
[42] KubeDL 2025. Run in Host Network. https://kubedl.io/docs/training/hostnetowrk/ Accessed on May 26, 2025.

[43] KubeLinter. 2024. KubeLinter: Static analysis for Kubernetes. https://github.com/stackrox/kube-linter/. Accessed on September 25, 2024.

[44] Kubernetes. 2024. Kubernetes Documentation: Headless Services. https://kubernetes.io/docs/concepts/services-networking/service/#headless-services. Accessed: December 2, 2024.

[45] Kubernetes. 2024. Kubernetes Documentation: Labels and Selectors. https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/. Accessed: December 2, 2024.

[46] Kubernetes. 2025. Kubernetes Documentation: Cluster Networking. https://kubernetes.io/docs/concepts/cluster-administration/networking/. Accessed: May 19, 2025.

[47] Kubesec 2024. Kubesec: Security risk analysis for Kubernetes resources. https://kubesec.io/ Accessed on September 25, 2024.

[48] Kyverno. 2024. Kyverno. https://kyverno.io/. Accessed on September 25, 2024.

[49] Wubin Li, Yves Lemieux, Jing Gao, Zhuofeng Zhao, and Yanbo Han. 2019. Service Mesh: Challenges, State of the Art, and Future Research Opportunities. In *2019 IEEE International Conference on Service-Oriented System Engineering (SOSE)*. 122–1225. https://doi.org/10.1109/SOSE.2019.00026

[50] Xing Li, Yan Chen, Zhiqiang Lin, Xiao Wang, and Jim Hao Chen. 2021. Automatic Policy Generation for Inter-Service Access Control of Microservices. In *USENIX Security 21*. 3971–3988.

[51] Guannan Liu, Xing Gao, Haining Wang, and Kun Sun. 2022. Exploring the unchartered space of container registry typosquatting. In *31st USENIX Security Symposium (USENIX Security 22)*. 35–51.

[52] Vijay B Mahajan and Sunil B Mane. 2022. Detection, Analysis and Countermeasures for Container based Misconfiguration using Docker and Kubernetes. In *2022 International Conference on Computing, Communication, Security and Intelligent Systems (IC3SIS)*. 1–6. https://doi.org/10.1109/IC3SIS54991.2022.9885293

[53] Microsoft. 2021. Secure containerized environments with updated threat matrix for Kubernetes. https://www.microsoft.com/en-us/security/blog/?p=93183 Accessed on September 25, 2024.

[54] Microsoft. 2024. Kubernetes Network Policies in Azure. https://learn.microsoft.com/en-us/azure/virtual-network/kubernetes-network-policies. Accessed on September 25, 2024.

[55] F. Minna, A. Blaise, F. Rebecchi, B. Chandrasekaran, and F. Massacci. 2021. Understanding the Security Implications of Kubernetes Networking. *IEEE Security & Privacy* 19, 05 (2021), 46–56. https://doi.org/10.1109/MSEC.2021.3094726

[56] Francesco Minna and Fabio Massacci. 2023. SoK: Run-time security for cloud microservices. Are we there yet? *Computers & Security* 127 (2023), 103119. https://doi.org/10.1016/j.cose.2023.103119

[57] MITRE 2024. MITRE ATT&CK. https://attack.mitre.org Accessed on September 25, 2024.

[58] Yuki Nakata, Katsuya Matsubara, and Ryosuke Matsumoto. 2021. Concentrated Isolation for Container Networks toward Application-Aware Sandbox Tailoring. In *Proceedings of the 14th IEEE/ACM International Conference on Utility and Cloud Computing (UCC '21)*. Article 15, 10 pages. https://doi.org/10.1145/3468737.3494092

[59] Jaehyun Nam, Seungsoo Lee, Hyunmin Seo, Phil Porras, Vinod Yegneswaran, and Seungwon Shin. 2020. BASTION: A Security Enforcement Network Stack for Container Networks. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*. USENIX Association, 81–95.

[60] Sam Newman. 2021. *Building microservices* (2nd ed.). O'Reilly Media.

[61] National Security Agency (NSA), Cybersecurity, and Infrastructure Security Agency (CISA). Aug. 2022. Kubernetes Hardening Guidance. https://www.nsa.gov/Press-Room/News-Highlights/Article/Article/2716980/nsa-cisa-release-kubernetes-hardening-guidance/.

[62] OWASP 2024. A05 Security Misconfiguration - OWASP Top 10:2021. https://owasp.org/Top10/A05_2021-Security_Misconfiguration/ Accessed on September 25, 2024.

[63] Prometheus Node Exporter 2025. Monitoring Linux host metrics with the Node Exporter. https://prometheus.io/docs/guides/node-exporter/ Accessed on May 26, 2025.

[64] Akond Rahman, Rezvan Mahdavi-Hezaveh, and Laurie Williams. 2019. A systematic mapping study of infrastructure as code research. *Information and Software Technology* 108 (2019), 65–77. https://doi.org/10.1016/j.infsof.2018.12.004

[65] Akond Rahman, Shazibul Islam Shamim, Dibyendu Brinto Bose, and Rahul Pandita. 2023. Security Misconfigurations in Open Source Kubernetes Manifests: An Empirical Study. *ACM Trans. Softw. Eng. Methodol.* 32, 4, Article 99 (may 2023), 36 pages. https://doi.org/10.1145/3579639

[66] Razi Rais, Christina Morillo, Evan Gilman, and Doug Barth. 2024. *Zero Trust Networks: Building Secure Systems in Untrusted Network* (2nd ed.). O'Reilly Media, Incorporated.

[67] Red Hat Inc. 2024. The state of Kubernetes security report: 2024 edition. https://www.redhat.com/en/resources/state-kubernetes-security-report-2024 Accessed on June 2, 2025.

[68] Liz Rice. 2020. *Container Security: Fundamental Technology Concepts that Protect Containerized Applications* (1st ed.). O'Reilly Media.

[69] Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly. 2020. *Zero trust architecture.* NIST Special Publication 800-207. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207

[70] Md Shazibul Islam Shamim, Farzana Ahamed Bhuiyan, and Akond Rahman. 2020. XI commandments of Kubernetes security: A systematization of knowledge related to Kubernetes security practices. In *2020 IEEE Secure Development.* 58–64. https://doi.org/10.1109/SecDev45635.2020.00025

[71] Shopify 2024. Kubeaudit: command line tool to audit Kubernetes clusters. https://github.com/Shopify/kubeaudit Accessed on September 25, 2024.

[72] Murugiah Souppaya, John Morello, and Karen Scarfone. 2017. *Application Container Security Guide.* NIST Special Publication 800-190. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-190

[73] StackRox. 2024. StackRox: Security platform for cloud-native applications, containers, serverless and Kubernetes. https://github.com/stackrox/stackrox Accessed on September 25, 2024.

[74] William Stallings. 2015. *Foundations of Modern Networking: SDN, NFV, QoE, IoT, and Cloud* (first ed.). Pearson. https://www.pearson.com/en-us/pearsonplus/p/9780137582235.html

[75] SUSE. 2024. NeuVector: Full Lifecycle Container Security Platform. https://www.suse.com/products/neuvector/ Accessed on September 25, 2024.

[76] Terraform Linters. 2024. tflint GitHub Repository. https://github.com/terraform-linters/tflint. Accessed on September 25, 2024.

[77] Thanos 2024. Thanos: Highly available Prometheus setup with long term storage capabilities. https://thanos.io/. Accessed: December 2, 2024.

[78] VMware. 2023. VMware Blog. https://tanzu.vmware.com/content/ebooks/stateofkubernetes-2023 Accessed on September 25, 2024.

[79] Nanzi Yang, Wenbo Shen, Jinku Li, Xunqi Liu, Xin Guo, and Jianfeng Ma. 2023. Take Over the Whole Cluster: Attacking Kubernetes via Excessive Permissions of Third-party Applications. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23).* 3048–3062. https://doi.org/10.1145/3576915.3623121

[80] Zegl 2024. Kube-score: Kubernetes object analysis with recommendations for improved reliability and security. https://github.com/zegl/kube-score Accessed on September 25, 2024.

[81] Hui Zhu and Christian Gehrmann. 2022. Kub-Sec, an automatic Kubernetes cluster AppArmor profile generation engine. In *2022 14th International Conference on COMmunication Systems & NETworkS (COMSNETS).* 129–137. https://doi.org/10.1109/COMSNETS53615.2022.9668504

A Disclosure

This section details the methodology employed for the responsible disclosure as well as its impact.

A.1 Methodology

We first identified the preferred private disclosure mechanism for each organization; if not specified, we contacted the maintainer of the chart. The disclosure included: a list of identified misconfigurations and the affected charts; the considered threat model (in Section 3.1) and a description of each type of misconfiguration, including suggestions on mitigations. We then engaged in an active and constructive discussion with these organizations so as to assess their severity through expert knowledge in the domain of the application. Specifically, we asked to internally circulate an anonymous questionnaire regarding our methodology and the severity of the misconfigurations. The questions (in Figure 5) concerned the role and experience of the respondent, considerations on security handling (including the use of security software), and feedback on the reported misconfigurations.

(1) **What is the size of your organization, if applicable? (number of employees)** [*Options: 1-99; 100-999; 1,000-4,999; 5000 or more; Not applicable*]
(2) **What is your current role?** Please describe your job position (e.g., Software Developer, SRE, DevOps Engineer) [*Text*]
(3) **How long have you been using Helm?** [*Options: Less than a year; 1-2 years; More than 2 years]*
(4) **Do you follow any guidelines to secure Helm Charts? If so, what are the main steps?** [*Text*]
(5) **Do you use any software tools or services to check the security of Helm Charts? If so, which?** [*Text*]
(6) **Compared to Charts created by your organization, do you handle third-party Helm Charts differently? (e.g., sanity checks, security checks, testing)** Third-party Charts are Charts created outside your organization. [*Text*]
(7) **What do you think of the following statements:** Lateral movement in Kubernetes means that an attacker gets control of other pods after getting a foothold into the cluster. Port information in a Helm chart defines the network ports used by the application's services. This includes specifying port names, port numbers, and target port numbers.
   - Detecting lateral movement in a Kubernetes cluster is a critical issue *[Options on a 5-point Likert scale]*
   - I trust the port information in Helm Charts *[Options on a 5-point Likert scale]*
(8) **Do you use network policies with your cloud applications?**[*Options: Yes; No*]
(9) **Why do you use network policies? What are their advantages and disadvantages?**[*Text, only if Yes was answered*]
(10) **Why you don't use network policies? What are their disadvantages?**[*Text, only if No was answered*]
(11) **Evaluate your agreement or disagreement with the following statements using the scale provided. Choose the response that best reflects your opinion on the following misconfigurations:** Undeclared ports are ports used by the container running in a pod, but not declared by the Chart / Pod. Unused ports are ports declared by the Chart / Pod but not used by the container.Label collision happens when different Kubernetes components use the same set of labels.
   - Undeclared ports are critical security risk *[Options on a 5-point Likert scale]*
   - Unused ports are a critical security risk *[Options on a 5-point Likert scale]*
   - Label collision is a critical security risk *[Options on a 5-point Likert scale]*
(12) **Why they are not a critical security risk?**[*Text, only shown if one of the previous answers to the critical risk was negative*]
(13) **Did you receive a security report about Helm misconfigurations, including Undeclared ports, Unused ports and / or Label collision?** Undeclared ports are ports used by the container running in a pod, but not declared by the Chart / Pod. Unused ports are ports declared by the Chart / Pod but not used by the container. Label collision happens when different Kubernetes components use the same set of labels. [*Options: Yes; No*]
(14) **Are there false positives in the reported misconfigurations?** False positives are components with unused / undeclared ports, or components with the exact same set of labels, but as a result of a design choice. [*Text*]
(15) **Evaluate your agreement or disagreement with the following statements using the scale provided. Choose the response that best reflects your opinion regarding the mitigation and detection of misconfigurations:**
   - The proposed mitigations are useful *[Options on a 5-point Likert scale]*
   - I will use a tool to detect the reported misconfigurations *[Options on a 5-point Likert scale]*
(16) **If the proposed mitigations were not useful, what would be a better option?** [*Text*]
(17) **Does the report reflect the status of your project? Leave here your feedback about the report** [*Text*]
(18) **Please leave here any other feedback you may consider useful for our research** [*Text*]

Fig. 5. Content of the feedback questionnaire used to follow-up to the disclosure.

A.2 Impact

We finally explain the impact of disclosing the misconfigurations we have found to the pertinent organizations at the time of writing. In particular, we show that the reported misconfigurations have been taken into serious consideration and also fixed in many cases.

Figure 6 illustrates that 22 pull requests have been created to fix the misconfigurations reported to Bitnami. In particular, Figure 8 shows a pull request for the RabbitMQ application – a widely-used message broker – to fix the *M1* misconfiguration (i.e., open ports are not declared).

Figure 7 shows several commits fixing misconfigurations in different Helm Charts of the European Environment Agency (EEA). Figure 9 shows one of them, adding a missing port to the Helm chart.

The disclosure to the Wikimedia foundation resulted in four patches to fix some port declaration issues and the removal of unused functionality from other charts. Figure 10 is an example of the changes on one application, which listened on an incorrect IP address.

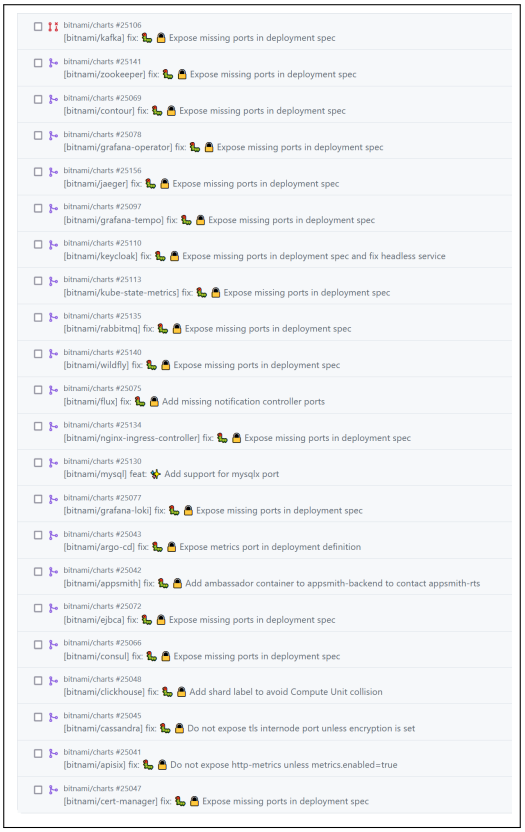

Fig. 6. Pull requests fixing misconfigurations in applications belonging to the Bitnami dataset.

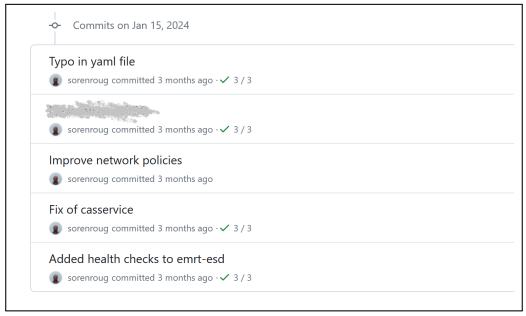

Fig. 7. Commits fixing misconfigurations in different charts belonging to the European Environment Agency (EEA) dataset.

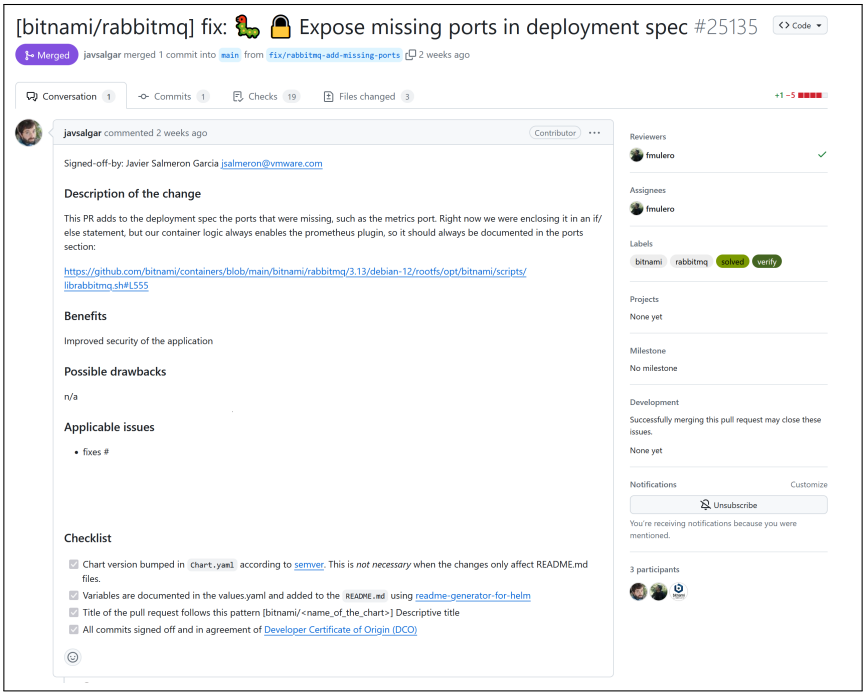

Fig. 8. Pull request fixing a misconfiguration in the `rabbitmq` chart of the Bitnami dataset.

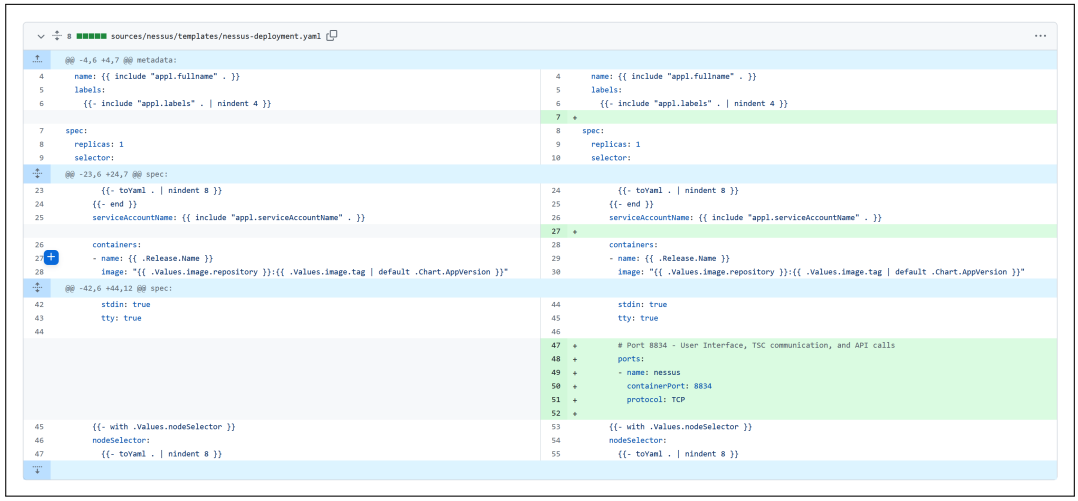

Fig. 9. Pull request fixing a misconfiguration in the nessus chart of the EEA dataset.

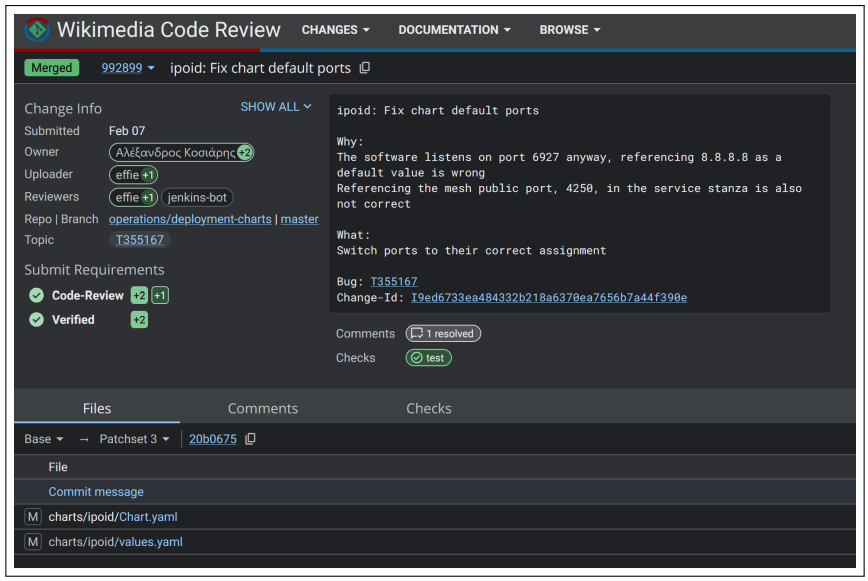

Fig. 10. Pull request fixing a misconfiguration in the `ipoid` chart of the Wikimedia foundation dataset.