An Empirical Study of Security Practices for Microservices Systems

Ali Rezaei Nasab, Mojtaba Shahin, Seyed Ali Hoseyni Raviz, Peng Liang, Amir Mashmool, and Valentina Lenarduzzi

Abstract—Despite the numerous benefits of microservices systems, security has been a critical issue in such systems. Several factors explain this difficulty, including a knowledge gap among microservices practitioners on properly securing a microservices system. To (partially) bridge this gap, we conducted an empirical study to manually analyze 861 security points collected from 10 GitHub open-source microservices systems and Stack Overflow posts concerning security of microservices systems, leading to a catalog of 28 microservices security practices. We then ran a survey with 63 microservices practitioners to evaluate the usefulness of these 28 practices. Our findings demonstrate that the survey respondents affirmed the usefulness of the 28 practices. These 28 security practices are further classified into six categories based on their topics: *Authorization and Authentication*, *Token and Credentials*, *Internal and External Microservices*, *Microservices Communications*, *Private Microservices*, and *Database and Environments*. We believe that the catalog of microservices security practices can serve as a valuable resource for microservices practitioners to more effectively address security issues in microservices systems. It can also inform the research community of the required or less explored areas to develop microservices-specific security practices and tools.

Index Terms—Microservices, Security, Empirical Study, Practitioners, Practice.

1 Introduction

Over the past few years, the Microservices Architecture (MSA) style has been popularly and widely used in the software industry. The MSA style aims to decompose a software application into or build it consisting of a set of microservices (i.e., small business-driven services) that can be implemented, tested, and deployed independently [1], [2]. Another benefit of this style is that it allows software development organizations to use the best-fit programming language and technology (e.g., database) to implement each microservice. Several other merits of microservices systems (software systems that employ the MSA style) are pointed out in the literature, such as scalability, modularity, and fault-tolerant [1], [3].

The research and industry communities have investigated several aspects of microservices systems. One of the most prevailing investigated and demanding aspects is how to migrate a legacy/monolithic application to microservices [4], [5]. For example, Balalaie et al. [6] identified 15 patterns for this purpose, and Auer et al. [7] developed an assessment framework to help software organizations specify and measure possible advantages and difficulties of the migration. Some researchers also looked at approaches and tools for microservices systems. Cinque et al. [8] developed an approach for monitoring microservices, and Heorhiadi et al. [9] introduced a framework, *Gremlin*, to test the "failure-handling capabilities" of microservices. Other researchers empirically investigated how microservices systems are designed, implemented, tested, monitored in the software industry (e.g., [10], [11]).

Since the advent of the MSA style, securing microservices systems has been a challenge for software practitioners and organizations [4], [12], [13], [14], [15], [16], [17]. The potential security challenges associated with microservices systems may compel software organizations to revisit their decision to adopt or migrate to microservices [18], [19]. The difficulty in securing microservices systems lies in several factors: (i) Tools and technologies that microservices use or rely on are prone to several security weaknesses and vulnerabilities; (ii) There is a knowledge gap among practitioners and organizations on securing microservices systems as the MSA style is an emerging and evolving architecture style; (iii) The distributed nature and characteristics of microservices systems make security harder than monolithic systems. For example, it is more difficult to guarantee security in such systems than monolithic systems as hundreds of microservices might be simultaneously running in production. Some works have been conducted on security in microservices systems (e.g., [12], [16], [20]), and recent review studies have called for more studies on security in microservices systems [1], [5], [15]. While these works

- *Ali Rezaei Nasab is with the School of Computer Science, Wuhan University, Wuhan, China, 430072. E-mail: rezaei.ali.nasab@gmail.com*
- *Mojtaba Shahin is with the Department of Software Systems and Cybersecurity, Faculty of IT, Monash University, Melbourne, Australia, 3800. E-mail: mojtaba.shahin@monash.edu*
- *Seyed Ali Hoseyni Raviz is with the School of Computer Science, Wuhan University, Wuhan, China, 430072. E-mail: s.ali.hoseyni@gmail.com*
- *Peng Liang is with the School of Computer Science, Wuhan University, Wuhan, China, 430072. E-mail: liangp@whu.edu.cn*
- *Amir Mashmool is with the Department of Computer Engineering, Faculty of Engineering, University of Birjand, Birjand, Iran, 9717434765. E-mail: amir.mashmool@birjand.ac.ir*
- *Valentina Lenarduzzi is with the Faculty of Information Technology and Electrical Engineering, University of Oulu, Oulu, Finland, FI-90014. E-mail: valentina.lenarduzzi@oulu.fi*
- *Peng Liang is the corresponding author.*

Manuscript received; revised.

are valuable for security in microservices systems, there is a lack of empirical evidence on the best microservices-specific security practices employed in the software industry and how practitioners perceive the usefulness of these best security practices. A comprehensive understanding of the best microservices-specific security practices and software practitioners' views on these practices would provide tailored support for designing and implementing secure microservices systems.

To address this gap, we conducted an empirical study that identified and validated 28 security practices for microservices systems. More specifically, we collected and analyzed 861 microservices security points gathered from 10 GitHub open-source microservices systems and Stack Overflow posts concerning security in microservices systems to collect the 28 security practices. These 861 security points are 543 GitHub issues, nine official documents, three wiki pages of the 10 microservices systems, and 306 Stack Overflow posts. We then ran an online validation survey completed by 63 microservices practitioners to evaluate the usefulness of the 28 security practices.

The key contributions of this paper can be summarized as follow:

- Identification of 28 security practices in six categories for microservices systems;
- Validation of the usefulness of these security practices from 63 microservices practitioners;
- Providing an online replication package of the data used in this study for researchers and practitioners to replicate and validate the findings [21];
- A set of actionable recommendations for microservices practitioners and researchers.

The rest of the paper is organized as follows: Section 2 provides the background on microservices systems and their security and summarizes the related work. Section 3 details our research methodology. The findings are presented in Section 4, followed by a set of recommendations for practitioners and researchers in Section 5. Section 6 elaborates on the threats of our study. Section 7 concludes our work and provides future work directions.

2 Background and Related Work

2.1 Background

2.1.1 Microservices Systems

Despite the MSA style being originated from Service-Oriented Architecture (SOA), they have some significant differences [1], [2]. Both include (small) services with dedicated responsibilities, while services in the MSA style are independent and autonomous and communicate through different lightweight mechanisms, services in SOA are not full-stack and fully autonomous [22].

The development, test, and deployment of each microservice can be done independently by a different development team using divergent technologies and programming languages. The unique characteristics of the MSA style allow microservices to scale independently from each other [3]. Furthermore, the MSA style enables the development teams to use the hardware that adequately meets their needs to deploy each microservice. As microservices are small and independent, their maintenance and making them fault-tolerant would be much easier. It is because the failure of one service will not lead to the entire system down, which may occur in monoliths [2].

However, the migration to or adopting microservices can be associated with many serious challenges and issues, which need careful consideration [23], [24]. These issues and challenges are enormous, such as the cost overhead due to the migration, the higher complexity of the system due to the increased amount and variety of components integrated into the system, and security issues [18], [19].

2.1.2 Security in Microservices Systems

Besides the various advantages brought by microservices, security becomes an issue when adopting and developing microservices systems. A monolithic system is a single system, and usually, a single application server needs to be secured, while in microservices systems, each microservice represents a possible attack surface [13].

In monolithic systems, communications between different components happen locally, with local calls, while in microservices systems, the communication happens through the network, creating another possible attack surface. In the MSA style, a compromised microservice can send malicious requests to other microservices. Another aspect to consider is the authorization between services since not all the services might be authorized to connect to other services. Different systems, such as Kubernetes[1] or Istio[2] provide inter-service authorization mechanisms. However, these mechanisms need to be developed, and new distributed access rules need to be defined separately for each service. The authorization mechanisms should try to reduce the privileges between services as much as possible instead of opening granting full access to all the services [13].

When dealing with authentications, distributed authentication mechanisms need to be considered. As an example, developers need to decide how to handle authentication, if using an authentication server, or independent authentication systems in different microservices [13]. Moreover, developers need to discriminate how to update the authentication mechanisms every time when new services or new users are included in the system.

On top of the security and authentication issues that need to be carefully designed by the architects, who are designing a specific system, microservices systems also need to consider possible vulnerabilities due to the usage of public container images that might be potentially infected [25]. An attempt to reduce this issue is provided by the "Moving Target Defenses", which proposes to modify component images to create uncertainty for attackers [26].

Moreover, when migrating from a monolithic system to microservices, companies need to keep the monolithic system and the microservices alive and connected at the same time until the migration is completed and the monolithic system is shut-down [23], [24]. This requires creating a secure communication channel between the monolithic system and the microservices system and integrating the

1. https://kubernetes.io
2. https://istio.io

security and authentication system adopted for the monolithic system with the new one adopted in the microservices system [23], [24]. The main challenges are to identify a suitable approach to:

- Manage and synchronize the authentication of microservices with the monolithic system, so as that end-users will not realize that the system is being migrated;
- Secure the microservices, the communication between the monolithic system and the microservices, but also the communication intra-microservices ;
- Understand which type of vulnerabilities are introduced during the migration.

2.2 Related Work

Several secondary studies have been conducted on (security in) microservices systems. In a recent review study [5], Waseem et al. observed that there are a few concrete solutions for addressing security concerns when implementing microservices systems in DevOps. Another systematic mapping review on 46 papes by Hannousse and Yahiouche [17] revealed that most studies on security in microservices are the solutions proposed in the soft-infrastructure layer. They argued that internal attacks, compared to external attacks, are less explored in the literature. They also indicated more efforts should be allocated on other layers of MSA (e.g., communication and deployment layers) and developing mitigation techniques. Pereira-Vale et al. [13] carried out a multivocal literature review on 36 academic literature and 34 grey literature. Their review led to a classification of 15 security mechanisms, in which authentication and authorization are the security mechanisms most reported in the literature. In addition, "mitigate/stop attacks" are expressed in about 2/3 of the security mechanisms. Ponce et al. [27] focused on security smells and refactorings and collected and analyzed 58 white and grey literature published from 2014 to 2021. They found ten security smells with their security properties and corresponding refactorings. They also elaborated on how the corresponding refactorings mitigate the effects of the smells.

Waseem et al. [14] empirically investigated the issues reported in five open source microservices systems hosted on GitHub. They found that 10.18% of these issues can be attributed to security. Yarygina and Bagge [16] asserted that security in microservices systems is a multi-faceted problem, and a layered security solution can address it. Hence, they categorized microservices security concerns into six layers (hardware, virtualization, cloud, communication, service/application, and orchestration) with solutions to address them (e.g., secure implementation of service discovery and registry components).

To secure IoT microservices, Pahl et al. [20] proposed a graph-based access control module in a network of IoT nodes. This module can monitor the communication of IoT microservices systems to create robust security and mitigate the security holes of such systems. Moreover, Pahl and Donini [28] provided a method based on "X.509 certificates" for authenticating IoT microservices. One of the main goals of this method is to verify the security properties locally through the distributed IoT nodes. Yu et al. [29] focused on security issues of microservices-based fog applications because such systems are composed of numerous microservices with complex communications, which is a challenge in terms of security. They reviewed 66 articles and identified 17 security issues (e.g., kernel exploit or DOS attack) regarding the microservices communication categorized in four groups: containers issues, data issues, permission issues, and network issues.

By surveying 67 participants, Rezaei Nasab et al. [12] affirmed that security challenges in microservices systems differ from non-microservices systems. To bridge the security knowledge gap among microservices practitioners, they developed a set of machine and deep learning approaches to automatically recognize security discussions (including security design decisions, challenges, or solutions) from open source microservices systems. Chondamrongkul et al. [30] also developed an automated approach using ontology reasoning and model checking techniques to identify security threats of microservices architectures through analyzing security characteristics. The identified security threats show how the attack scenarios may happen. Sun et al. [31] designed an API primitive FlowTap that provides security-as-a-service for microservices-based cloud applications. The proposed technique can monitor and protect the network of such systems from internal and external threats.

In contrast to the works above, our study presents 28 best security practices for microservices systems. These practices were collected from developer discussions occurring during the development of microservices systems. We further validated the usefulness of these practices by seeking feedback from 63 microservices practitioners. Finally, we articulated the positive and negative sides (if any) of the 28 practices.

3 Methodology

Our goal in this study is to identify and evaluate a catalog of security practices for microservice systems. Two research questions (RQs) are formulated to reach this goal.

> ***RQ1.*** *What are the security practices to secure microservices systems?*

Motivation. Documented knowledge and guidelines on how software practitioners design and implement secure microservices systems are scarce (if any) [12]. This RQ aims to fill this gap by developing a catalog of best security practices employed by practitioners to secure microservices systems. To this end, we manually analyzed developer discussions in GitHub open-source microservices systems hosted and Stack Overflow posts concerning security in microservices systems (see Section 3.1).

> ***RQ2.*** *To what extent do practitioners consider the identified microservices security practices useful?*

Motivation. The security practices identified from GitHub and Stack Overflow are based on the authors' analysis, which might be subjective and unreliable. This RQ aims to

evaluate and generalize the security practices identified in **RQ1** by conducting an online survey. The survey seeks the perception of microservices practitioners about the usefulness of the identified security practices (see Section 3.2).

3.1 Mining Security Practices (RQ1)

We describe how we collected developer discussions related to microservices security in Section 3.1.1. We then report how we analyzed the collected data to identify microservices security practices in Section 3.1.2.

3.1.1 Data Collection

Developer discussions reported on GitHub and Q&A websites (e.g., Stack Overflow and Stack Exchange) are valuable sources for identifying development practices (e.g., architectural practices [32], [33], security practices [34]). Considering this, we tried to build a dataset of developer discussions about security in microservices systems with the following steps.

Step 1. In our previous work [12], we manually created a dataset of 5,018 security paragraphs collected from 567 issues in seven open-source microservices systems and 505 posts from Stack Overflow. These security paragraphs include *"design decisions, challenges, or solutions relating to security in microservices systems"* [12]. The architectural style of these seven open-source systems (i.e., goa, eShopOnContainers, microservices-demo, scalecube-services, moleculer, deep-framework, and light-4j shown in Table 2) was determined as the microservices architecture style by enquiring the core contributors of those projects through an online survey. The core contributors are defined as *"the top three contributors who have the most commits in a project"* [12]. Our previous work [12] aimed to develop ML/DL-based approaches to automatically differentiate security paragraphs from non-security paragraphs in microservices developer discussions. In contrast, this work focuses on manually analyzing developer discussions on security to identify and evaluate best security practices for microservices systems.

Hence, the 5,018 security paragraphs collected from our previous work could be a good source to identify security practices. Despite this, the security paragraphs are short (i.e., 2-3 sentences), and there might be the chance of losing the design context by reading the security paragraphs individually. Our strategy to (partially) mitigate this challenge was to read the entire issues or posts that entail equal or more than five such security paragraphs. Our decision to set the threshold to five security paragraphs was based on our experience in the previous work [12] as issues and posts with five or more security paragraphs are more suitable for identifying security practices. As shown in Figure 1, this process led to 76 issues and 306 posts, which are called security points. The creation date of these 306 Stack Overflow posts ranges from August 2014 to July 2020.

Step 2. Apart from the seven open-source microservices systems used in our previous work [12], we found three larger open-source microservices systems (spinnaker with 6,514 issues, jaeger with 3,222 issues, and cortex with 4,438 issues shown in Table 2) on GitHub. These three new systems also employ the MSA style. We used the same approach in [12] (i.e., enquiring the core contributors) to identify these new microservices systems.

We applied DeepM1 (i.e., the best performing ML/DL approach developed in [12]) to identify security paragraphs from these three new projects. Since the unit of analysis in the approach of [12] was a paragraph, we converted the issues of these three projects to paragraphs using their HTML tag <p>. The pre-processing used in [12] was also used for these paragraphs. Note that DeepM1 with a recall of 84.25% and a precision of 86.73% works at the paragraph level. This process led to identifying 9,566 security paragraphs from these three new systems, which include 5,931 security paragraphs from spinnaker project, 1,629 security paragraphs from jaeger project, and 2,006 security paragraphs from cortex project (see Figure 1). Similar to the approach used in **Step 1**, we selected issues that entail equal or more than five such security paragraphs. This resulted in a collection of 467 security points from spinnaker, cortex, and jaeger (see Figure 1).

Step 3. Among all systems discussed in **Step 1** and **Step 2**, nine systems (i.e., eShopOnContainers, spinnaker, cortex, jaeger, goa, moleculer, deep-framework, microservices-demo, and light-4j) have documentation with security discussions. Further, three systems, including eShopOnContainers, scalecube-services, and light-4j, have a wiki page with discussions on security. We considered these nine pieces of documentation and three wiki pages as security points. As shown in Figure 1, our dataset includes 861 security points, which were manually analyzed to identify security practices (see Section 3.1.2).

3.1.2 Identification of Security Practices

Identifying and extracting microservices security practices from the 861 microservices security points include three phases.

Pilot Phase. We randomly selected 20 security points from the 861 security points and asked three analysts (three authors) to analyze them independently. The goal was to get familiar with data and understand what sort of developer discussions should be considered a security practice. Each analyst applied the open coding and constant comparison techniques from Grounded Theory [35] to extract security practices. They then held a meeting to check the similarities and dissimilarities of the extracted security practices and resolve any disagreements.

Main Phase. The three analysts collaboratively analyzed the remaining 841 security points. 90 security points were allocated to the three analysts each week. In other words, each analyst was asked to extract security practices from 30 allocated security points using the open coding and constant comparison techniques [35]. Further, an Excel file was created and shared with all the analysts. They were requested to maintain the link between an identified practice and its corresponding security point. At the end of each week, the analysts held a meeting to discuss their extracted security practices, identify the duplicate ones, merge, or rephrase them. This process led to the identification of 36 practices.

Feedback Phase. In this phase, the other three authors reviewed the 36 security practices. They mainly checked the 36 security practices to identify and mitigate possible inconsistencies and ambiguities. This step resulted in merging four practices with other microservices security

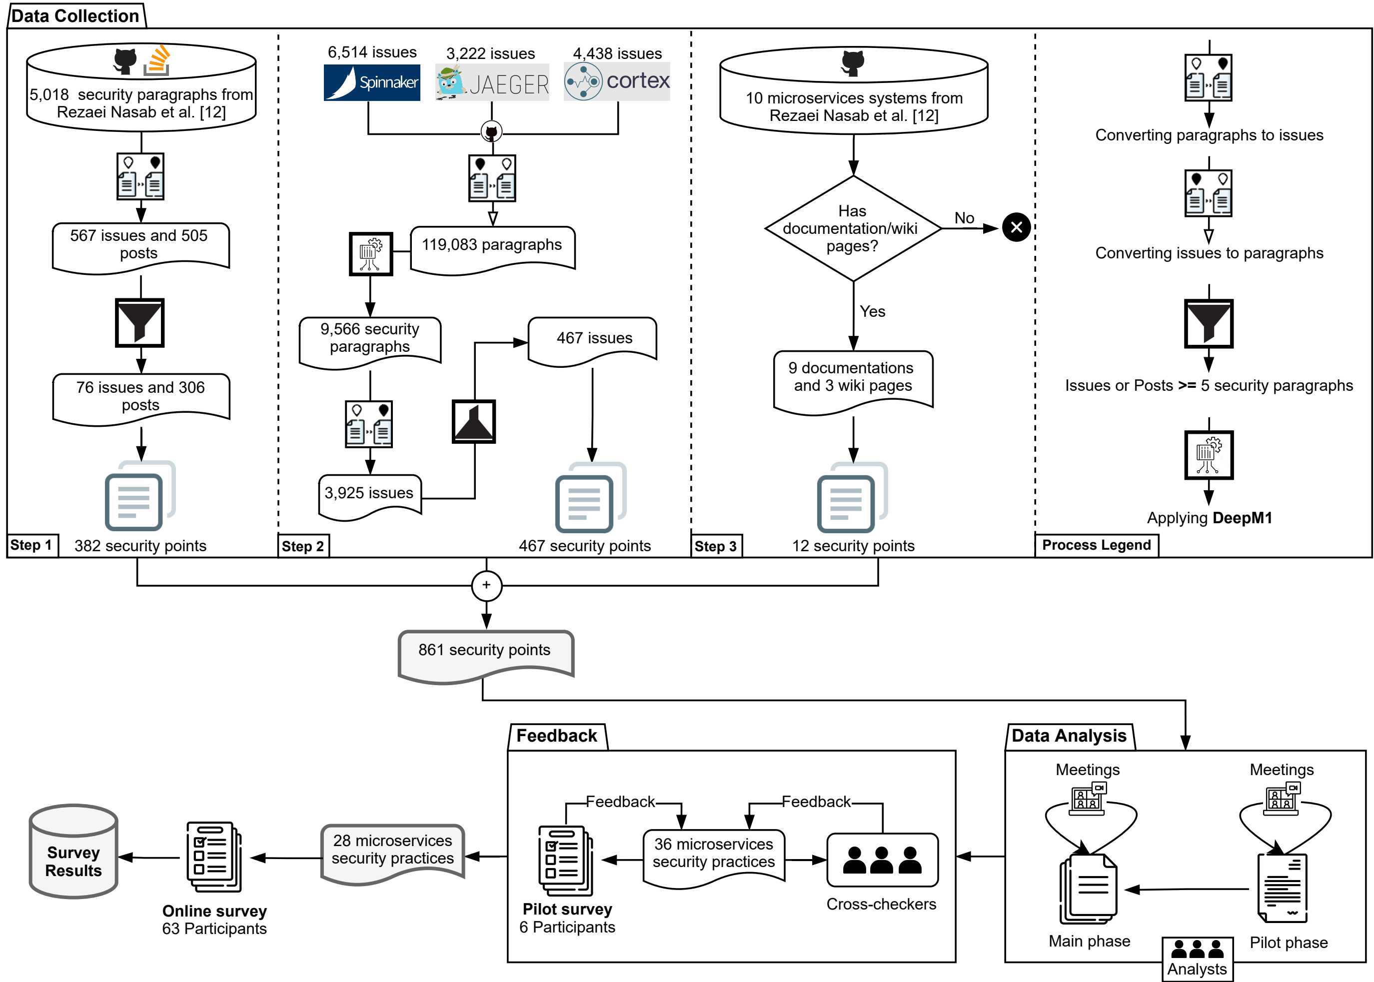

Fig. 1. An overview of extracting microservices security practices

TABLE 1
Demographic questions of our online survey

Demographic Questions	Question Type	Example Answers
"How many years have you been involved in software development?"	Multiple choice	0 <year <= 1, 1 <year <= 2
"What is your main role in software development?"	Multiple choice	Developer, Architect, Tester
"How many years have you been involved in microservices system development?"	Multiple choice	0 <year <= 1, 1 <year <= 2
"How many years of experience do you have with security in microservices systems?"	Multiple choice	0 <year <= 1, 1 <year <= 2
"How large is your organization?"	Multiple choice	20 <= employees <= 50
"What are the domains of your organization?"	Checkbox	Financial, E-commerce
"Which country do you currently work in?"	Free Text	Australia

practices and reducing the number of security practices to 32 practices.

Next, we ran a pilot survey to seek microservices practitioners' feedback on the 32 identified security practices. The pilot survey was designed using Google Forms and completed by six microservices practitioners. We asked the practitioners to demonstrate their level of agreement or disagreement with the 32 identified security practices (Likert scale questions rated from "strongly agree = 5" to "strongly disagree = 1"). Each Likert scale question was followed by an optional open-ended question to obtain further feedback. Based on feedback collected from the practitioners and our internal discussions, we reduced the number of security practices from 32 to 28. This reduction was because the practitioners indicated that two security practices did not contain enough information to be understood. We also found two practices had overlapped and merged them. We also slightly rephrased the wording of some practices (e.g., adding more information to a security practice) to remove any ambiguities. We classified these 28 security practices into six categories based on their topics: *Authorization and Authentication* (6 practices), *Token and Credentials* (5 practices), *Internal and External Microservices* (7 practices), *Microservices Communications* (4 practices), *Private Microservices* (2 practices), and *Database and Environments* (4 practices).

3.2 Validation Survey (RQ2)

In this section, we elaborate on the design and execution of the validation survey.

3.2.1 Protocol

Considering the guidelines proposed by Kitchenham and Pfleege [36], an online survey (i.e., the validation survey) was developed to evaluate the usefulness of the 28 microservices security practices identified in Section 3.1. The survey was anonymous and hosted on Google Forms. The survey preamble describes the goal of the survey and briefly explains how and from which sources these 28 microservices security practices are identified. The survey includes 45 questions and takes about 25 minutes to complete. The survey questions can be classified into three groups.

Demographic questions. We asked seven questions to get background information about the survey participants (e.g., *"how many years of experience do you have with security in microservices systems?"*). Table 1 shows these seven questions. All demographic questions except one were compulsory.

Likert scale questions. The respondents were asked to rate the usefulness of each of the 28 identified security practices using a Likert question. The mandatory Likert scale questions were ranked on a four-point scale as *"Absolutely Useful"*=4, *"Useful"*=3, *"Not Useful"*=2, and *"Absolutely Not Useful"*=1. We also added the option *"I Don't Know"* to allow practitioners not to respond to the practices when they were unsure about or unclear to them.

Open-ended questions. As discussed in Section 3.1.2, the 28 security practices are classified into six categories. For each category, we asked the participants to provide the reason for their response for one of the practices in that category that they rated *"Absolutely Useful/Useful"* or *"Absolutely Not Useful/Not Useful"*. The participants were also requested to list the practice number. Note that answering these questions was optional. Finally, two more optional questions were asked. The respondents were requested to share any feedback about the security practices in microservices systems. The second question was to allow the participants to provide their email addresses if they were interested in the results of our study.

3.2.2 Participants

We used the following methods to recruit microservices practitioners.

❶ We collected the publicly available emails of 868 contributors involved in the ten projects listed in Table 2. We emailed them and asked them to fill up the survey.

❷ The spinnaker project has a workplace on Slack with many active contributors. The workplace has some Special Interest Group (SIG) channels that focus on different topics (e.g., security). We advertised our survey in the workplace.

❸ The third strategy was to broadly advertise the survey in some microservices groups on social networks like Twitter and LinkedIn. Further, we sent private messages to practitioners who were a member of these groups.

❹ We asked the invited practitioners to share the survey with their colleagues who had experience in microservices security.

In total, we received 63 valid responses. The initial analysis of the survey responses revealed that two responses were invalid. For example, one participant answered all 28 Likert questions as *"I Don't Know"*. Note that we did not calculate the response rate for our survey because of our heterogeneous recruitment process (e.g., the respondents might contribute to one or more projects in Table 2, and at the same time, they might be involved in multiple LinkedIn groups).

3.2.3 Data Analysis

Descriptive statistics were used to study the responses to the closed-ended questions, i.e., demographic and Likert scale questions. We also applied the open coding technique to analyze the responses to the open-ended questions [35]. Note that we used the answers (if any) to the open-ended questions to clarify why a particular security practice was chosen *"Useful/Absolutely Useful"* or *"Not Useful/Absolutely Not Useful"* by the respondents.

4 Findings

4.1 Demographics

Here we provide the background information about 63 practitioners who responded to the survey.

Experience. Figure 2 (A) shows that 50.8% of the participants have been involved in software development for at least 10 years. All participants, except for one participant, had at least one year of experience in developing microservices systems, with 50.8% having more than three years of experience (see Figure 2 (B)). Regarding Figure 2 (C), more than 70% of the respondents had one year of experience with securing microservices systems. 23.8% worked with security in microservices systems in less than one year. The rest (4.8%) did not have any experience in this regard.

Role. As shown in Figure 2 (D), the participants mainly worked as Developer (30.2%, 19 out of 63), Software Engineer (25.4%, 16 out of 63), Architect (23.8%, 15 out of 63), and Technical Lead (12.6%, 8 out of 63).

Organization size and domain. The majority of the participants (85.6%) came from organizations with more than 100 employees (see Figure 2 (E)). 60.3% (38 out of 63 participants) were from organizations with more than 1000 employees. The participants' organization domains are shown in Figure 3. The participants were able to choose one or more organization domains in the demographic question. The dominant domains are financial and consulting and IT services, followed by E-commerce and telecommunications.

Country. The distribution of participants per country is shown in Figure 4. Since this question was optional, we only received 42 responses for this question. The 42 participants who indicated their country information came from 22 countries across six continents, including Europe (10 countries), Asia (5 countries), North America (2 countries), South America (2 countries), Africa (2 countries), and Oceania (1 country). Most of them were from India, Brazil, and Australia (see Figure 4).

TABLE 2
A list of ten microservices systems used in this study. Release (Rel.); Contributors (Cont.); Line of Codes (LoC)

#	Project Name	URL	Stars	Forks	Issues	Rel.	Cont.	LoC	Languages	Docs	Wiki
1	eShopOnContainers	https://bit.ly/2X40b4M	18.6k	7.9k	1,752	17	142	120k	C#	✔	✔
2	jaeger	https://bit.ly/2YvyJgN	14.2k	1.7k	3,222	44	221	105k	Go, Shell	✔	✘
3	spinnaker	https://bit.ly/3ndjJyu	8k	1.1k	6,514	132	116	6k	Shell, Go	✔	✘
4	moleculer	https://bit.ly/2X5DayD	4.6k	441	1,003	99	92	98k	Javascript	✔	✘
5	goa	https://bit.ly/38LUYkD	4.4k	459	2,907	55	88	88k	Go	✔	✘
6	cortex	https://bit.ly/3nbq65x	4.3k	604	4,438	50	204	1.2m	Go	✔	✘
7	light-4j	https://bit.ly/3zOnhL0	3.3k	554	1,030	140	34	57k	Java	✔	✔
8	microservices-demo	https://bit.ly/38LtxY2	2.8k	1.8k	875	13	55	15k	Python	✔	✘
9	deep-framework	https://bit.ly/3ncbgM4	538	75	647	22	12	956k	Javascript	✔	✘
10	scalecube-services	https://bit.ly/3DTFhGn	507	79	820	178	21	11k	Java	✘	✔

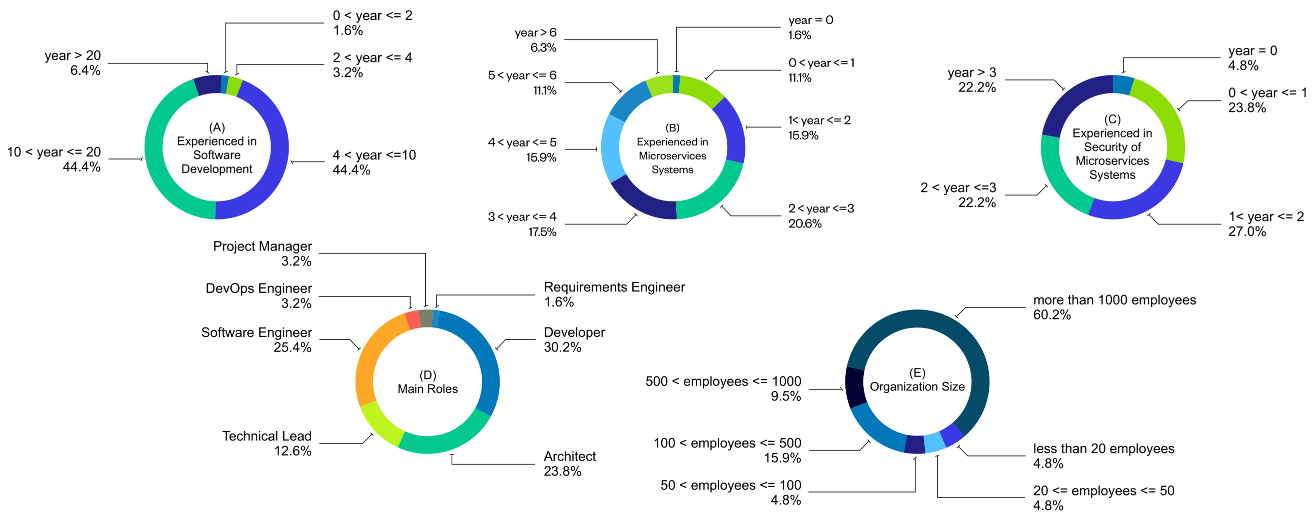

Fig. 2. Experience of the participants (n=63) in software development (**A**), experience of the participants in microservices system development (**B**), experience of the participants in securing microservices systems (**C**), main roles of the participants (**D**), organization size of the participants (**E**).

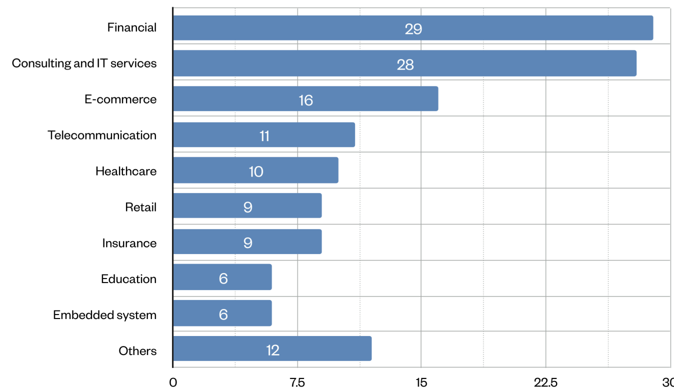

Fig. 3. Participants' organization domains (n=63). Note: participants could select more than one domain

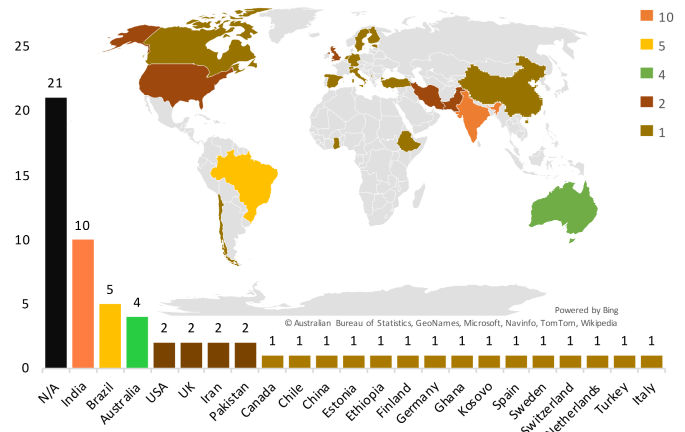

Fig. 4. Number of participants (n=63) from countries

4.2 Security Practices

This section details the 28 security practices identified through mining 861 security points (RQ1). We also present the perspective of the survey respondents about the usefulness of these security practices (RQ2). As described in Section 3.1.2, these 28 practices are classified into six groups.

Similar to the arguments by Malavolta et al. [33], the main goal behind revealing the usefulness of these security practices is to indicate their applicability in practice and assess the reliability of our analysis of the practices. We do not aim to rank these practices. Still, practitioners need to consider the design context and the requirements and constraints of

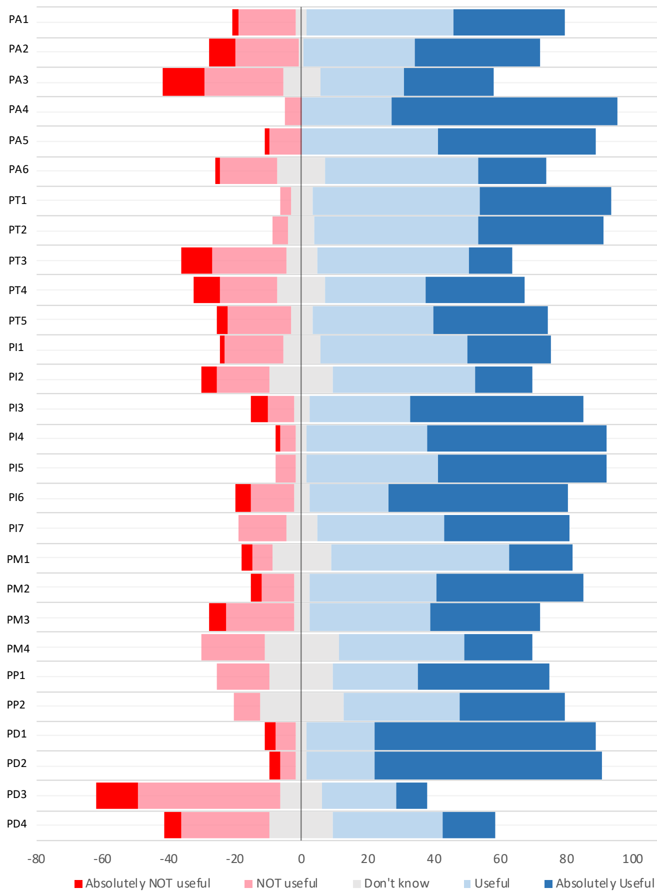

Fig. 5. The level of usefulness of the 28 identified security practices for microservices systems from the perspective of 63 practitioners

their microservices systems and organizations when using these practices (see Section 5.1).

4.2.1 Authorization and Authentication

This group includes six practices for authorizing microservices or authenticating users in microservices systems (see Table 3).

PA1. *Add an "identity microservice", and authorize microservices access through "identity microservice" and token of each microservice. Each microservice has its own token, which is passed in each request between the microservices.* A microservices system includes several microservices, and each of the microservices performs specific tasks. Mostly, they need to communicate with each other to reach out to a target. Hence, authorizing microservices is an important task in these systems. Assume a user wants to be authenticated once and then access all the relevant microservices and client apps with their protected resources. This process is referred to as Single Sign-On (SSO) [13]. **PA1** is a way to implement SSO [37]. In this practice, a microservice called "identity microservice" will check the token of microservices related to the user and make sure they are valid microservices. In case they are not valid microservices, the user will be redirected for the authentication.

This practice may negatively affect *performance* due to too many round trips to the authentication server. Also, it affects *scalability* because of the increased number of microservices. 49 (77.7%) out of the 63 survey respondents confirmed that this practice is *(absolutely) useful*. Some of the comments that the participants posted to support or refuse this practice are:

> *"[It is] useful to **maintain traceability** in requests between microservices."* (Architect)

> *"I believe that microservices should have one service that is dedicated to authentication. Authentication service may be developed inside or subscribed as a service from other companies. Default gateway routes all incoming requests to the authentication service then the **authentication service validates whether the request has a valid token or not.** In my experience, each microservice has a private key of JWT, which checks all incoming requests before allowing access to a resource."* (Software Engineer)

> *"[I am] not sure [about] the **latency impact** [of this practice]."* (DevOps Engineer)

> *"When working with microservices, it is important to remember that you have to follow your good practices, if you pass the token to all microservices, they will have to call the microservice for authorization always, in this case it is **faster and safer to use middleware** in the API gateway, this way centralizes authentication."* (Software Engineer)

PA2. *Each microservice in the microservices architecture must be responsible for its own security, i.e., each microservice must have security enabled.* Another way to implement SSO is to use a decentralized authentication protocol like OpenID, with which each microservice can handle its own security [37]. In other words, with this practice, the user will have to authorize each microservice individually. More than half of the respondents (71.5%) acknowledged that the practice is *absolutely useful* or *useful*. The following are two examples of the participants' comments that support **PA2**.

> *"Each microservice should enable security and have to **check all incoming requests before allowing access to the resource**, but [other stuff related to] microservice security such as token creation, refresh token, OAuth implementation should be handled by a separately dedicated authentication service."* (Software Engineer)

> *"Each service needs to be **protected with security token**."* (Developer)

On the other hand, some respondents mentioned that it is not necessary to secure all microservices:

> *"**[It is] not necessarily all microservices** in the architecture should be concerned with security, as this should be in the gateway API or in the infrastructure layer."* (Software Engineer)

PA3. *Suppose you use an API Gateway approach in the microservices architecture. In that case, you do not need to have each microservice security enabled (you do not need to implement* **PA2***) because the internal microservices can be protected by not being published out of the Docker host.* One of the key challenges for practitioners is to decide to authorize microservices at the microservices level only, authorize at the API Gateway level, or both of them [48]. Suppose they choose the API Gateway approach for this purpose. In that case, the microservices could only be accessed by other Containers within the Docker host through the "internal port" of each Container. 29 survey participants (52.4%) considered this practice *(absolutely) useful*. On the other hand, 21 (36.5%)

TABLE 3
Security practices for **authorizing and authentication** and the survey responses (in %). **AU**: Absolutely Useful, **U**: Useful, **NU**: Not Useful, **ANU**: Absolutely Useful, **IDK**: I Don't Know, **MED**: Median, **AVG**: Average

		Sources	AU	U	NU	ANU	IDK	MED	AVG
PA1	Add an "identity microservice", and authorize microservices access through "identity microservice" and token of each microservice. Each microservice has its own token, which is passed in each request between the microservices.	[37], [38], [39], [40], [41], [42], [43], [44]	33.3	44.4	17.5	1.6	3.2	3	3.13
PA2	Each microservice in the microservices architecture must be responsible for its own security, i.e., each microservice must have security enabled.	[37], [45], [46], [47]	38.2	33.3	19.0	7.9	1.6	3	3.03
PA3	Suppose you use an API Gateway approach in the microservices architecture. In that case, you do not need to have each microservice security enabled (you do not need to implement **PA2**) because the internal microservices can be protected by not being published out of the Docker Host.	[48]	27.0	25.4	23.8	12.7	11.1	3	2.75
PA4	A large microservices system is recommended to use an API Gateway approach for securing/authorizing/routing microservices.	[48], [47], [49], [50], [51]	68.4	27.0	4.8	0.0	0.0	4	3.63
PA5	Set an authorization boundary on each microservice even when (1) microservices are "internal" and (2) it is possible to set the authorization at the API Gateway level.	[52]	47.6	41.3	9.5	1.6	0.0	3	3.39
PA6	Use Public Key Infrastructure (PKI) signing/verification system to prevent round trips to the Authorization Service.	[53]	20.6	46.0	17.5	1.6	14.3	3	3

opposed it (23.8% rated it as *not useful*, and 12.7% rated it as *absolutely not useful*).

Some participants indicated that authorizing microservices at the microservices level and the API Gateway level is required for microservices systems as there should be zero trust security in such systems.

👎 *"You can't know who is making the requests and you should have authentication/authorization enabled (**zero trust**)"* (Software Engineer)

👎 *"Having an API gateway does not mean to disable security in microservices. In fact API gateway is used for **coarse grained security issues**, microservices are used for **fine grained security**. Both levels of security are required."* (Requirements Engineer)

💡 **PA4.** *A large microservices system is recommended to use an API Gateway approach for securing/authorizing/routing microservices.* The API Gateway approach causes overhead in small microservices systems and needs too many steps when coding and updating features [48], [47]. Out of twenty-eight practices, **PA4** was mostly rated by the survey participants as *absolutely useful* or *useful* (more than 95%). No one voted this practice as *absolutely not useful*.

We received only positive comments on this practice, indicating that it works well for routing, minimizes coupling, and supports the evolution of microservices.

👍 *"I agree default **gateway [can be used] for routing** but authentication and authorization should be handled by separate service."* (Software Engineer)

👍 *"Using gateway makes it easier for the client because (1) they don't have to deal with many different addresses of each service, (2) abstraction is highly enforced which also enhances security because **endpoints of the services are not directly exposed**, thereby making it difficult for attackers, and (3) **coupling is minimised**."* (Developer)

👍 *"We are using it, and **it enables routing correctly**."* (Technical Lead)

💡 **PA5.** *Set an authorization boundary on each microservice even when (1) microservices are "internal" and (2) it is possible to set the authorization at the API Gateway.* This practice is recommended when developers want to add authorization with Ocelot[3] [52]. However, developers need to balance between security and simplicity. The majority of our survey respondents agreed with the usefulness of this practice (88.9% rated it as *absolutely useful* or *useful*).

As shown in the following comments, **PA5** (1) provides a native cloud solution, (2) is useful for accidental wrong access/modification, and (3) reduces the security concerns in internal microservices.

👍 *"It uses a **native cloud solution**."* (DevOps Engineer)

👍 *"Authorisation boundary is very important and useful not only in terms of security but also for **accidental wrong access/modification**."* (Software Engineer)

👍 *"[It] is cool because you create an **outer layer of security** and your internal services don't have to worry about security."* (Software Engineer)

In contrast, a detractor mentioned that:

👎 *"Some requests navigate through many services before returning a response, so implementing authorization in API Gateway, I believe, is not the best practice. The default gateway should only do the **routing stuff**, adding some security-related task to the default gateway is not a **scalable solution**. Because if you want to implement SSO in the future, I think it's a bit **difficult to implement** it."* (Software Engineer)

💡 **PA6.** *Use Public Key Infrastructure (PKI) signing/verification system to prevent round trips to the Authorization Service.* This practice aims to handle the authorization in a microservices environment [53]. Most respondents (66.6% considered this practice *useful* or *absolutely useful*. As an example of the positive comments from the respondents on this practice, we have:

👍 *"If we use PKI then **each microservice can validate the security tokens** instead of sending a request to identity microservice to validate the tokens."* (Developer)

Detractors noted that the use of PKI in a microservices architecture is costly.

👎 *"Not useful, depending on the number of requests between microservices, performing key verification **can be very costly**."* (Architect)

3. https://ocelot.readthedocs.io/en/latest/index.html#

*"This will have a **high cost of development and maintenance**."* (DevOps Engineer)

4.2.2 Token and Credentials

As shown in Table 4, this group includes five practices for handling sensitive information in a microservices system.

PT1. *Use a method based on the Public/Private key to secure microservices through JSON Web Token (JWT).* JWTs can be signed in a microservices system and generate two key pairs (private signing key and public verification key) [64]. The public verification key generated by a JWT can be distributed to all microservices in the microservices system. If microservice A wants to decrypt the information in microservice B, it only needs to know the private signing key created by microservice B [54]. If the microservices system uses the API Gateway approach, the API Gateway should also know the private signing key.

90.5% (57) of the survey respondents rated it as *absolutely useful* or *useful*. Only two practitioners chose *not useful*. This practice received only positive comments, in which the respondents offered to use the OAuth stream in addition to JWT.

*"JWT alone is not enough. [It is] interesting to use an **OAuth stream** with client credentials + cookie, in a Gateway API strategy."* (Architect)

*"For external facing APIs, we can use **OAuth based authentication**."* (DevOps Engineer)

PT2. *In microservices systems, use JSON Web Tokens (JWTs) to handle the session expiration/revocation.* This practice recommends using Redis tool[4] to track token revocations [37]. **PT2** received almost similar positive feedback to **PT1** from the survey respondents (87.3% *absolutely useful* or *useful*). A respondent stated that this practice is useful if someone uses JWT for communication between microservices behind the gateway. Another participant confirmed **PT2** as a useful practice, but he/she mentioned that JWT is not the only method to handle the session expiration and revocation.

*"Communication between **microservices behind the gateway** can be **JWT** which is a value token. But from **client to gateway** should be a **reference token** which does not contain any sensitive information."* (Architect)

"JWT is not the only option." (Software Engineer)

PT3. *Secure caches of credentials in each microservice that needs to access other microservices.* More than 55% of the practitioners verified that **PT3** [60] is *absolutely useful* or *useful*, while 31.7% rated it as *not useful* or *absolutely not useful*. A software engineer with more than three years of experience in security of microservices systems pointed out:

*"If the cached credentials **are not secured**, the entire credential system is **questionable**."* (Software Engineer)

Some respondents believed that (1) the usefulness of **PT3** depends on whether the microservice is stateful or stateless, and (2) the overhead of securing credentials.

*"It depends on the [micro]service type weather it is **stateful** or **stateless**. If the service is stateful, we may think about a way how to store the credential and use it in the next request. Else we use a private key to validate the request token."* (Software Engineer)

*"**Security is necessary but an overhead**, only secure what needs securing. Maintain separation of concerns, **manage/cache user credentials in one dedicated service**."* (Technical Lead)

PT4. *Decode JSON Web Token (JWT) at the microservices level instead of the API Gateway level.* 38 participants (60.4%) believed that JWTs should be decoded at the microservices level because they mostly include relevant information for authentication and authorization. API Gateway can manage the JWTs in the form of Fail-fast (a.k.a. fail early), and it is just recommended to verify access tokens at the microservices level [61]. Our analysis shows that 16 respondents (25.4%) did not agree with the usefulness of **PT4**. Below are two comments that question **PT4** and rationalize why decoding JWTs should be done at the API Gateway level.

*"Since JSON token is decoded once at the gateway level, **the overhead of each service having to deal with the decoding is removed**, thereby making availability better."* (Developer)

*"I believe decoding can be done **at the API Gateway level** and from there, the request should opt for a different way to hit the internal services. That will free the services from doing any kind of decoding work. I think it would have **better performance and be more secure**."* (Architect)

PT5. *Endpoints of microservices like server information, health check, and logging level must be secured in the request/response chain.* Our analysis shows that server information, health check, and logging level contain sensitive information and should be secured as part of securing a microservices system [62], [63]. Depending on the requirements of a microservices system, developers may only use some of these endpoints. 71.5% of the respondents opted for *absolutely useful* or *useful* for this practice. Some comments that indicate the importance of securing endpoints are shown as:

*"Leaving **diagnostics information** unsecured may expose loopholes in the system, which makes it easier for attackers. Again sensitive information can be leaked to unauthorized users."* (Developer)

*"All endpoints, including **diagnostic endpoints, must be secured**. These are prone to attacks and can leak potentially sensitive data."* (Technical Lead)

*"Specially **logs could contain data** that need to be protected at all time."* (DevOps Engineer)

In contrast, a respondent disagreed with protecting health checks as it may not allow the implementation of a fault tolerance strategy.

*"**Health checks** should **not** be protected, in a fault tolerance strategy whoever makes the request needs to know if the microservice is active, to allow a retry alternative."* (Architect)

4.2.3 Internal and External Microservices

This group of practices focuses on securing a set of microservices. Part of these microservices (internal microservices) is used inside an organization, and the rest (external microservices) may be used by any third-party (see Table 5).

4. https://redis.io

TABLE 4
Security practices for **tokens and credentials** and the survey responses (in %). **AU**: Absolutely Useful, **U**: Useful, **NU**: Not Useful, **ANU**: Absolutely Useful, **IDK**: I Don't Know, **MED**: Median, **AVG**: Average

		Sources	AU	U	NU	ANU	IDK	MED	AVG
PT1	Use a method based on the Public/Private key to secure microservices through JSON Web Token (JWT).	[54], [55], [56], [57], [58]	39.7	50.8	3.2	0.0	6.3	3	3.39
PT2	In microservices systems, use JSON Web Tokens (JWTs) to handle the session expiration/revocation.	[37], [59]	38.1	49.2	4.8	0.0	7.9	3	3.36
PT3	Secure caches of credentials in each microservice that needs to access other microservices.	[60]	12.8	46.0	22.2	9.5	9.5	3	2.68
PT4	Decode JSON Web Token (JWT) at the microservices level instead of the API Gateway level.	[61]	30.2	30.2	17.5	7.9	14.2	3	2.96
PT5	Endpoints of microservices like server information, health check, and logging level must be secured in the request/response chain.	[62], [63], [52]	34.9	36.6	19.0	3.2	6.3	3	3.1

TABLE 5
Security practices for **internal and external microservices** and the survey responses (in %). **AU**: Absolutely Useful, **U**: Useful, **NU**: Not Useful, **ANU**: Absolutely Useful, **IDK**: I Don't Know, **MED**: Median, **AVG**: Average

		Sources	AU	U	NU	ANU	IDK	MED	AVG
PI1	Developers can use internal microservices secured with a different token than external microservices. In this scenario, API Gateway acts as a token issuer for the internal microservices.	[65]	25.4	44.4	17.5	1.6	11.1	3	3.05
PI2	Developers can use internal microservices secured using the tokens of external microservices, and their permission must be controlled using Access Control List (ACL). In this scenario, API Gateway forwards the tokens to the internal microservices.	[65]	17.5	42.9	15.8	4.8	19.0	3	2.9
PI3	Whether microservices are only internally used within an organization or are externally accessible to third parties, authentication is required either way.	[66], [67], [68]	52.4	30.2	7.8	4.8	4.8	4	3.37
PI4	In an internal microservice use case, "client credential" should not get exposed to the third party.	[69]	54.0	36.5	4.7	1.6	3.2	4	3.48
PI5	Microservices systems made of components should be isolated and internal calls should not be leaked outside their boundaries.	[70]	50.8	39.7	6.3	0.0	3.2	4	3.46
PI6	Encrypt tokens if they are going to be exposed to the outside of the system boundary.	[71]	54.0	23.7	12.7	4.8	4.8	4	3.33
PI7	It is recommended to minimize the number of HTTP dependencies between internal microservices. This will minimize the future impact on microservices performance and Denial-of-Service attacks.	[72]	38.1	38.1	14.3	0.0	9.5	3	3.26

PI1. *Developers can use internal microservices secured with a different token than external microservices. In this scenario, API Gateway acts as a token issuer for the internal microservices.* 44 respondents (69.8%) acknowledged that developers could use unique tokens for securing internal and external microservices. One practitioner rated it *absolutely not useful* (1.6%), and seven practitioners considered it *not useful* (17.5%). Below are two negative comments on this practice.

> *"The default gateway should not handle* ***Identity and Access Management (IAM)*** *task."* (Software Engineer)

> *"There is* ***no need to use token and secure*** *internal microservices as long as they are not accessible from outside."* (Software Engineer)

PI2. *Developers can use internal microservices secured using the tokens of external microservices, and their permission must be controlled using Access Control List (ACL). In this scenario, API Gateway forwards the tokens to the internal microservices.* Similar to **PI1**, **PI2** aims to secure internal microservices. However, it uses the tokens of external microservices for securing internal microservices [65]. 60.4% of the participants stated that **PI2** is *absolutely useful* or *useful*. 20.6% rated this practice as *not useful* or *absolutely not useful*. The following comment includes negative feedback on this practice.

> *"I don't know how useful is to having the same token internally and externally and* ***encrypting the token is always the best****."* (Architect)

PI3. *Whether microservices are only internally used within an organization or are externally accessible to third parties, authentication is required either way.* The majority of the survey respondents (82.6%, 52 out of 63) considered **PI3** as *absolutely useful* or *useful*. They argued that the authenticated microservices remain secure in the following scenarios: (1) the occurrence of misconfigurations that lead to exposure of internal microservices to outside, and (2) if a security hole is opened in the firewall [66].

> *"If authentication is not enabled for internal microservices, then as soon as the* ***internal physical network gets compromised****, the entire microservices system is compromised, which is a disaster."* (Software Engineer)

On the other hand, some respondents believed that this practice would be only necessary or useful under certain circumstances.

> *"It* ***depends on the service*** *that the microservice gives. Some services may need authentication, and some may not need a user to authenticate."* (Developer)

"*By having communication via the [message] broker, it is not necessary to authenticate in internal microservices, but* ***if you use REST in microservices****, this is making a bad practice, and in that case,* ***it will be necessary****.*" (DevOps Engineer)

PI4. *In an internal microservice use case, "client credential" should not get exposed to the third party.* The client credentials are identifiers for accessing client data. It is strongly advised to distinguish internal client credentials from external ones [69]. This practice was rated as *absolutely useful* or *useful* by more than 90% participants. Only one respondent noted that if the third party is valid for the internal microservices, there is no problem exposing client credentials to the third party.

"*If we are talking about the grant type client credentials in OAuth2 and the* ***client ID/secret identifies the third party*** *system, I don't see a problem exchanging the "client credential" with the third party.*" (Requirements Engineer)

PI5. *Microservices systems made of components should be isolated and internal calls should not be leaked outside their boundaries.* This practice has more focus on controlling the components' exposure. Most participants (91.5%, 57 out of 63) marked it as *absolutely useful* or *useful*. Similar to **PI4**, none of the participants rated this practice as *absolutely not useful*.

PI6. *Encrypt tokens if they are going to be exposed to the outside of the system boundary.* From the perspective of a software developer with 10 years of experience, it is needed to encrypt the tokens because they contain some authorization-related information [71]. This practice was *absolutely useful* for 34 respondents (54%) and *useful* for 15 respondents (23.7%). The following are two positive comments which state that the tokens should be encrypted at all times.

"*Tokens can* ***be hacked*** *so should* ***always*** *be encrypted.*" (Architect)

"*Tokens should* ***be encrypted with a public key*** *and get decrypted with a private key regardless of the fact that the internal service is receiving it or a client.*" (Architect)

However, a survey respondent questioned the need for adding another layer of encryption if tokens are already signed.

"*Assuming tokens are* ***already signed****, what is the* ***reason*** *to have* ***another layer*** *of encryption?*" (Software Engineer)

PI7. *It is recommended to minimize the number of HTTP dependencies between internal microservices. This will minimize the future impact on microservices performance and Denial-of-Service attacks.* Our analysis of the security points revealed that some developers advised that fewer communications between internal microservices are better because being autonomous and available to the client is one of the purposes of microservices. If we employ HTTP dependencies between microservices, it can violate the autonomy of microservices [72]. It also impacts the performance of microservices when one of them does not perform well [72].

48 (76.2%) participants considered **PI7** as *absolutely useful* or *useful*. The survey respondents pointed out that it would be useful to avoid HTTP calls as much as possible because they may create some problems for internal microservices calls. Furthermore, the respondents emphasized that HTTP dependencies should be reduced because it is against the separation of concerns principle in the design of microservices systems. They also recommended using gRPC instead of HTTP for synchronous calls. An alternative to prevent the Distributed Denial-of-Service (DDoS) attack is to use the Backends For Frontends (BFF) pattern[5].

"*HTTP calls are synchronous. Hence too much use of it for inter-service calls may* ***cause availability issues****. If indeed synchronous calls are required, then* ***gRPC*** *may be used.*" (Developer)

"*One way to prevent* ***DDoS*** *is to work with* ***Backends For Frontends (BFF)*** *and only expose it to the world, and not expose each of your microservices to be accessible by external HTTP requests.*" (Software Engineer)

"***Separation of concerns*** *is a major feature of pure microservices.*" (Software Engineer)

A negative comment that we received on this practice is:

"*When we came to microservice architecture,* ***the most used way of synchronous messaging between services was by using HTTP*** *so as many requests could be sent and received to do the task. I don't recommend minimizing the number of requests as a solution.*" (Developer)

4.2.4 Microservices Communications

Table 6 represents four practices that are related to authenticating and authorizing requests when two or more microservices are communicating.

PM1. *Use OAuth2 "Client Credentials Flow" if two microservices that trust each other want to talk together from the backends.* In such communications, there is no end-user identity involved. **PM1** emphasizes the trust between microservices where they explicitly call each other [73]. Assume that there are a user and two microservices (A, B). The user accesses microservice A through a JSON Web Token (JWT). At this time, microservice A needs to access microservice B. The "OAuth2 client credentials grant" is recommended to handle the communication between these two microservices. If two microservices do not trust each other, the "OAuth2 client credentials grant" provides a good way to handle the authentication between these two microservices [74]. In this case, each microservice will use its own credentials to obtain a token through the "token microservice" (i.e., a microservice that is responsible for generating, renewing, and validating a token) and use it to connect to another microservice.

A large number of the survey participants (i.e., more than 70%) considered this as an (*absolutely*) *useful* practice. A few were the opposite of it (less than 10%). As a positive comment on this practice, we have:

"*Client credential is only valid when the* ***intercommunication*** *does not specify the* ***current active user****.*" (Architect)

We only received a negative comment on this practice.

"*Intercommunications between microservices* ***should be a custom grant type****, not client credential.*" (Architect)

PM2. *The connection between a microservice and its respective database should be protected by a security protocol, like Transport Layer Security (TLS).* Our analysis of the collected

5. https://samnewman.io/patterns/architectural/bff

TABLE 6
Security practices for **microservices communications** and the survey responses (in %). **AU**: Absolutely Useful, **U**: Useful, **NU**: Not Useful, **ANU**: Absolutely Useful, **IDK**: I Don't Know, **MED**: Median, **AVG**: Average

		Sources	AU	U	NU	ANU	IDK	MED	AVG
PM1	Use OAuth2 "Client Credentials Flow" if two microservices that trust each other want to talk together from the backends.	[73], [74], [75], [76]	19.0	54.0	6.3	3.2	17.5	3	3.08
PM2	The connection between a microservice and its respective database should be protected by a security protocol, like Transport Layer Security (TLS).	[77], [78]	44.4	38.1	9.5	3.2	4.8	3	3.3
PM3	When a microservice in a microservices architecture needs to call another microservice, the access token should be passed around microservices with the request.	[79]	33.3	36.5	20.6	4.8	4.8	3	3.03
PM4	It is recommended to use the gRPC framework for internal microservice-to-microservice synchronous communication.	[80], [81], [76], [82], [83], [84]	20.6	38.2	19.0	0.0	22.2	3	3.02

security points taught us that developers should be worried about the security of communication between a microservice and its database (or even other databases) [77]. As an essential part of the data protection strategy, Microsoft strongly advises protecting data in transit [85]. Moreover, because the data is exchanged from many locations, SSL or TLS protocols are highly recommended. 52 respondents admitted this practice (44.4% *absolutely useful* and 38.1% *useful*).

They mentioned that **PM2** is useful to prevent unauthorized access of microservices to the database. They also recommended that using the gRPC framework with the TLS protocol is a good practice when a microservice wants to communicate to its own database.

"It is better to secure the connection of a service and its database because it ***prevents unauthorized access*** *to the database."* (Developer)

"It's good to use ***gRPC framework and TLS*** *while accessing the DB."* (Architect)

Some others argued that (1) there is no need to use **PM2** once the database is in a private network, and (2) adding a security protocol in a connection causes more complexity.

"The DB must ***be embedded*** *in the service container or in a private network. Therefore there is* ***no need to encrypt*** *a local connection."* (Architect)

"Adding SSL to database connection only ***adds more complexity****."* (Software Engineer)

PM3. *When a microservice in a microservices architecture needs to call another microservice, the access token should be passed around microservices with the request.* Imagine there are a user and two microservices (A and B). The user and microservice A are in Scope A (i.e., the user is authorized for microservice A). Microservice B is in Scope B (i.e., the user is not authorized for microservice B). Suppose the user with the access token wants to use a resource from microservice A and at the same time, microservice A must call microservice B to give the resource to the user. In that case, **PM3** is recommended to prevent any communication failure [79].

69.8% of the proponents agreed **PM3** is *(absolutely) useful*. A participant provided a condition for the usefulness of this practice in the following comment.

"Only if you are going to get some information from the user, otherwise, it [passing the access token around microservices with the request] is not necessary." (DevOps Engineer)

PM4. *It is recommended to use the gRPC framework for internal microservice-to-microservice synchronous communication.* gRPC is a communication protocol using HTTP2 [86] and Protocol Buffers [87]. The analysis of the security points indicated that gRPC provides an effective solution for direct synchronous communication between microservices [80], [81]. 37 respondents agreed that **PM4** is a (*absolutely*) *useful* practice (58.8%). More than 20% of the participants were not familar with this practice. No one selected *absolutely not useful* for this practice.

Proponents argued that gRPC is faster than HTTP because it uses binary encoding. However, the costs of using it should be considered.

"gRPC uses binary encoding, which makes it ***faster****."* (Developer)

"You can take ***benefit over HTTP****. But it depends."* (Technical Lead)

"gRPC is useful but it has a ***cost****, you should think about its* ***advantages to implement*** *it."* (Software Engineer)

In contrast, a few respondents pointed out that **PM4** is not useful and should be avoided because, e.g., it causes difficulties in development and debugging.

"This is an option, but should ***not be a requirement****."* (Software Engineer)

"Binary data transfer makes development and debugging ***very difficult*** *and does* ***not add much of value*** *in terms of security or network performance."* (Software Engineer)

"gRPC should be ***avoided at any time****."* (Architect)

4.2.5 Private Microservices

Table 7 provides two practices to increase the security of private microservices. Private microservices are internal microservices in an organization that only a specific group of end-users or applications can access.

PP1. *Remote nodes (e.g., remote microservices) should not be able to even list/check for the existence of private microservices.* If there are some private microservices in a microservices architecture, none of the remote nodes must be allowed to call private microservices' actions [70]. In addition, they should not be allowed to check any information about private microservices directly (even the number of private microservices). In this scenario, private microservices can only contact internal microservices. 41 out of 63 survey respondents (65.3%) considered this practice *absolutely useful* or *useful*. We also did not receive *absolutely not useful* for this

TABLE 7
Security practices for **private microservices** and the survey responses (in %). **AU**: Absolutely Useful, **U**: Useful, **NU**: Not Useful, **ANU**: Absolutely Useful, **IDK**: I Don't Know, **MED**: Median, **AVG**: Average

		Sources	AU	U	NU	ANU	IDK	MED	AVG
PP1	Remote nodes (e.g., remote microservices) should not be able to even list/check for the existence of private microservices.	[70]	39.7	25.4	15.9	0.0	19.0	3	3.29
PP2	Use "service grouping" to limit the visibility and callability (of both actions and events) of the private microservices in a group of microservices.	[70]	31.8	34.9	7.9	0.0	25.4	3	3.32

practice. One participant shared the positive feedback on this practice as follow:

> *"The requirement/responsibility of checking for the existence [of private microservices] may cause **tight coupling between services** which may cause availability issues and also potential "distributed monolith" because a service may not function if other services are unavailable."* (Developer)

PP2. *Use "service grouping" to limit the visibility and callability (of both actions and events) of the private microservices in a group of microservices.* More than 65% of the survey participants (42 out of 63) acknowledged that the use of service grouping for private microservices is (*absolutely*) *useful*. This practice received the highest rate of "I don't know" among all practices (25.4%). However, **PP2** did not receive *absolutely not useful*, and only a few participants (7.9%) rated it as *not useful*.

An architect commented that service grouping is a beneficial method in microservices systems and is always required.

> *"As a best practice from Microservices or Containerization point of view, **service grouping is always required** and beneficial."* (Architect)

4.2.6 Database and Environments

Four practices are categorized into the database and environments group and are shown in Table 8. They focus on security concerns that databases and production environments may raise in microservices systems.

PD1. *Although security policies should be applied in both development and production environments, production environments need stronger security.* The majority of the survey participants (87.3%) accepted that production environments need more security policies than development environments. Among these participants, more than 65% rated it as *absolutely useful*. Less than 10% disagreed with the usefulness of this practice. In the following, we received some comments that support or refute/question **PD1**:

> *"Because different kinds of clients with different **unknown intentions will access** the system."* (Developer)

> *"It is best to keep development and production environments **as similar as possible** to avoid **surprise issues**."* (Software Engineer)

> *"Best practice would **enable the same** in development environment."* (DevOps Engineer)

> *"Depending on the information, the development environment needs to have a **security level equivalent to** the production one."* (Software Engineer)

> *"Security should be **repeated across all environments**."* (Software Engineer)

From the comments, some participants believed that the production environment should enable security policies more than the development one because various types of clients with unknown intentions can use the microservices system in the production environment. Conversely, other participants believed that development and production environments should have equal security policies to make them identical as much as possible.

PD2. *In a microservices architecture, databases should not be exposed to any unauthenticated request.* The clients send various requests to resources in a microservices architecture. Since the resources always contain important information, it is extremely recommended that no resources accept the requests which are not authenticated [89]. We received a high rate of usefulness (*absolutely useful* or *useful*) for **PD2** (88.9%) which more than 65% of them were *absolutely useful*. A few comments which support or refute this practice are shown below:

> *"Unauthenticated requests should not be enabled on database."* (Technical Lead)

> *"Needs to have its **default secure authentication**, as well as its own private VPN network."* (Software Engineer)

> *"There are some cases that don't require authentication and it requires to do database operation, e.g., **scheduler to clean up log table**."* (Architect)

PD3. *Suppose a microservice needs to validate some data against data from another microservice synchronously. In that case, it is recommended to combine both microservices and have only one microservice.* Almost 70% of the survey participants disagreed with the usefulness of the practice or indicated that they had no idea about this practice (see Table 8). Still, 31.7% of our survey participants chose **PD3** as (*absolutely*) *useful*.

The main reason stated this practice is not useful is that it can be against the Single Responsibility principle and increase the size of microservices.

> *"[In] some **exceptional cases,** this item might be useful but in most of the cases if we follow this we **will end up a few huge services** instead of a real microservices system."* (Software Engineer)

> *"To handle the synchronization scenario, I don't think that it is a good idea to **break the Single Responsibility principle** rather eventual consistency needs to follow."* (Architect)

PD4. *Suppose a microservice needs to validate some data against data from another microservice synchronously. In that case, the first microservice should replicate data from the second microservice in its own database with an eventual consistency syncing system.* Nearly 50% of the survey participants agreed

TABLE 8
Security practices for **database and environments** and the survey responses (in %). **AU**: Absolutely Useful, **U**: Useful, **NU**: Not Useful, **ANU**: Absolutely Useful, **IDK**: I Don't Know, **MED**: Median, **AVG**: Average

		Sources	AU	U	NU	ANU	IDK	MED	AVG
PD1	Although security policies should be applied in both development and production environments, production environments need stronger security.	[72], [88], [84]	66.7	20.6	6.3	3.2	3.2	4	3.56
PD2	In a microservices architecture, databases should not be exposed to any unauthenticated request.	[89]	68.3	20.6	4.7	3.2	3.2	4	3.59
PD3	Suppose a microservice needs to validate some data against data from another microservice synchronously. In that case, it is recommended to combine both microservices and have only one microservice.	[90]	9.5	22.2	42.9	12.7	12.7	2	2.33
PD4	Suppose a microservice needs to validate some data against data from another microservice synchronously. In that case, the first microservice should replicate data from the second microservice in its own database with an eventual consistency syncing system.	[90]	15.9	33.3	27.0	4.8	19.0	3	2.75

with this practice, and less than 35% of our survey participants voted **PD4** as (*absolutely*) *not useful*. One positive comment on this practice is:

👍 *"The recommendation in case of a sync would be to actually **keep a copy of the data** in the two services, remembering the **Consistency-Availability-Partition tolerance (CAP) theorem**."* (Software Engineer)

5 RECOMMENDATIONS

In this section, we present concrete and actionable recommendations for microservices practitioners and researchers based on our reflections on the findings.

5.1 Recommendation for Practitioners

The most useful practices. All 28 practices, except for practice **PD3**, tend to have the median Likert score of 3 or 4, indicating that the vast majority of the survey participants affirmed these practices are *useful* or *absolutely useful*. At the same time, we acknowledge that software organizations and practitioners may not be willing to or cannot adopt all 28 security practices (e.g., lack of enough resources). Hence, we highlight the seven most important security practices, including **PA4**, **PI3**, **PI4**, **PI5**, **PI6**, **PD1**, and **PD2**, with a median Likert score of 4 (*absolutely useful*) and encourage microservices practitioners to adopt these security practices to ensure the desired level of security in microservices systems.

PA4 from the "authorization and authentication" group emphasizes using API Gateway to handle authorizing or routing microservices in large-scale microservices systems. Four out of these seven highly affirmed practices, including **PI3**, **PI4**, **PI5**, and **PI6**, come from the "internal and external microservices" group. This can (partially) show the importance of the network of microservices in terms of internal and external domains and the level of security policies and practices that should be considered. **PD1** and **PD2** focus on security concerns in production environments and databases. **PD1** argues that microservices systems need more security policies when executing in production environments compared to development environments. **PD2** informs microservices practitioners that microservices' databases should not accept any types of unauthenticated requests.

Design context still matters. Although the survey participants considered almost all identified security practices useful, we do not claim that these 28 practices are the best options for all contexts and domains. We assert that practitioners should carefully consider different context-sensitive factors and trade-offs when using each of these practices. For example, practices **PA1** and **PA2** have the same goal to authorize microservices, and each can be used in a specific context to achieve the purpose. According to the responses of our participants, such factors and trade-offs are enormous, such as is the given system or service public or private? What are the impacts of the security practices on other quality attributes? How sensitive is the data? A Technical Lead with more than five years of experience in microservices systems commented on this as: "*All scenarios [practices] are useful, depending on system requirements. Security is an essential overhead; however, it is an overhead (with speed and complexity). So be sensible, don't over-engineer it. Equally, make sure sensitive data cannot be leaked - use private networks where possible.*"

5.2 Recommendations for Researchers

Study how the identified practices are used in different microservices systems in different domains and contexts. In Section 5.1, we discussed that although most of the survey participants affirmed the usefulness of the identified security practices, the successful implementation of these practices depends on too many factors and trade-offs. In this study, we have tried, to some extent, to show in which circumstances some practices are useful (e.g., **PT3**, **PI7**) or their impacts on other quality attributes (e.g., **PA1**). However, it is out of the scope of our study to explore and discuss all factors and trade-offs. For example, it is worth investigating the short-term and long-term impacts of each of these practices? What are their impacts on a specific type of microservices systems, e.g., IoT microservices systems? What are the associated costs and overhead of each of these practices? Further to this, there are still some practices that were slightly controversial (e.g., **PT4**, **PI3**, and **PD1**). For example, **PT4** suggests decoding JWT at the microservices level instead of the API Gateway level. 25.4% of the practitioners still thought that decoding JWT at the API Gateway level is better than at the microservices level. Finally, the survey respondents tended to disagree with a few security

practices (e.g., **PD3** and **PD4**) that we found from GitHub and Stack Overflow. We argue that future research should explore the abovementioned points through, e.g., controlled experiments, action research, case studies, or observational studies.

Pay attention to security practices in other or less explored aspects of microservices systems. In this study, we only looked at two developer discussion platforms (10 microservices systems from GitHub and 306 posts from Stack Overflow) to identify 28 security practices categorized into six groups. We believe that many security practices were not discussed or found in our data sources. While some other works (e.g., [13], [17]) have recently examined (gray) literature to understand security in microservices systems, we encourage researchers to mine other sources, such as other developer discussion platforms (e.g., Reddit[6]) to identify more practices. These new practices can either complement our security practices in a given group (e.g., the "database and environments" group) or be new categories of security practices (e.g., security practices for containerized microservices systems).

6 THREATS TO VALIDITY

In this section, we summarize potential threats to the validity of this study and the strategies that we used to mitigate these threats [91].

6.1 External Validity

Three threats might limit the generalizability of our findings. First, we identified the 28 security practices from only two data sources: GitHub and Stack Overflow. Although these platforms are the most popular online platforms among different types of software practitioners to share and discuss software development challenges, knowledge, and solutions, they do not represent all views of software practitioners. Second, we chose 10 open-source microservices systems on GitHub, which is only one popular OSS repository. These projects vary in terms of domain, number of contributors, size, etc. Despite this fact, we cannot claim that these projects are representative of all types of microservices systems (e.g., IoT microservices systems) and all the OSS repositories. Finally, our validation study did not receive a large number of responses (i.e., 63 responses). This threat was, to some extent, reduced as software practitioners with different characteristics (e.g., possessing different roles and working in organizations with diverse domains) completed the survey.

6.2 Internal Validity

Identifying security practices from Stack Overflow and GitHub might be subjective and error-prone. We adopted several strategies to reduce this issue. First, several analysts and validators were involved in this process. Three analysts participated in the pilot phase and the main phase of the data analysis (see Section 3.1.2). Three other authors with extensive experience in security in microservices systems reviewed and validated the identified security practices and suggested some feedback. Finally, we deployed a pilot survey to seek practitioners' feedback on the identified security practices. This helped us remove four practices and improved the wording of some of the security practices.

The validation survey tried to recruit practitioners with experience in securing microservices systems. Hence, we carefully checked the profiles of many practitioners on their websites, Slack, LinkedIn, etc. Still, practitioners with poor knowledge of the MSA style and security could concern the validity of the validation survey. We added the "I Don't Know" option in the survey questions to minimize this concern. We also asked the survey respondents to comment on why they rated a given practice "Useful" or "Not Useful". The detailed comments from the survey respondents increased our confidence that the vast majority of them had the right experience and expertise.

A survey's results might be biased if only practitioners with a few specific roles participate. Practitioners with different roles, such as developers, software engineers, DevOps engineers, requirements engineers, architects, completed the validation survey.

6.3 Construct Validity

Our decision to use DeepM1 introduced in [12] to extract security paragraphs from cortex, spinnaker, and jaeger projects might have introduced threats. Although DeepM1 has a good performance in detecting security paragraphs from developer discussions in microservices systems, we acknowledge that we might have missed some important security information from these projects. Furthermore, we defined security points as an issue or post with equal to or more than five security paragraphs. We admit that some issues or posts with less than five security paragraphs may still contain important microservices security practices. In this study, we only used the validation survey to evaluate the usefulness of the identified practices. Other research methods such as case studies could also be used to indicate all positive and negative aspects of the identified security practices.

6.4 Reliability

There is a potential threat that other researchers replicate our study and generate different results. Our first approach to alleviate this threat was to provide a detailed explanation of our research method (e.g., the survey design), enabling other researchers to replicate our study. Furthermore, we created a replication package [21], including the 861 security points used to identify security practices and the encoded survey responses, allowing other researchers and practitioners to validate our findings.

7 CONCLUSIONS AND FUTURE WORK

This study identified 28 security practices for securing microservices systems through manually examining 861 microservices security points. These 861 microservices security points include 543 GitHub issues, nine official documents, and three wiki pages from 10 open-source microservices systems, and 306 Stack Overflow posts concerning security in microservices systems. These 28 security practices are

6. https://www.reddit.com

categorized into six groups: *Authorization and Authentication*, *Token and Credentials*, *Internal and External Microservices*, *Microservices Communications*, *Private Microservices*, and *Database and Environments*. Through an online survey completed by 63 microservices practitioners, we have shown that the majority of the respondents rated these 28 practices useful for industrial usage.

In the future, we plan to extend our catalog of security practices by exploring more resources (e.g., interviews) to identify more security practices, in particular, in less explored areas of microservices systems. We also aim to investigate the positive and negative impacts of the identified security practices in different types of microservices systems.

Acknowledgments

This work was partially funded by the National Key R&D Program of China with Grant No. 2018YFB1402800 and the National Natural Science Foundation of China (NSFC) with Grant No. 62172311. The authors would also like to thank all the survey participants.

References

[1] N. Dragoni, S. Giallorenzo, A. L. Lafuente, M. Mazzara, F. Montesi, R. Mustafin, and L. Safina, "Microservices: yesterday, today, and tomorrow," in *Present and Ulterior Software Engineering*. Springer, 2017, pp. 195–216.

[2] M. Fowler and J. Lewis. (2014) Microservices a definition of this new architectural term. [Online]. Available: https://bit.ly/3zk5xXr

[3] P. Jamshidi, C. Pahl, N. C. Mendonça, J. Lewis, and S. Tilkov, "Microservices: The journey so far and challenges ahead," *IEEE Software*, vol. 35, no. 3, pp. 24–35, 2018.

[4] P. Di Francesco, P. Lago, and I. Malavolta, "Architecting with microservices: A systematic mapping study," *Journal of Systems and Software*, vol. 150, pp. 77–97, 2019.

[5] M. Waseem, P. Liang, and M. Shahin, "A systematic mapping study on microservices architecture in devops," *Journal of Systems and Software*, vol. 170, p. 110798, 2020.

[6] A. Balalaie, A. Heydarnoori, and P. Jamshidi, "Microservices architecture enables devops: Migration to a cloud-native architecture," *IEEE Software*, vol. 33, no. 3, pp. 42–52, 2016.

[7] F. Auer, V. Lenarduzzi, M. Felderer, and D. Taibi, "From monolithic systems to microservices: an assessment framework," *Information and Software Technology*, vol. 137, p. 106600, 2021.

[8] M. Cinque, R. Della Corte, and A. Pecchia, "Microservices monitoring with event logs and black box execution tracing," *IEEE Transactions on Services Computing*, 2019.

[9] V. Heorhiadi, S. Rajagopalan, H. Jamjoom, M. K. Reiter, and V. Sekar, "Gremlin: Systematic resilience testing of microservices," in *Proceedings of the 36th IEEE International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2016, pp. 57–66.

[10] M. Waseem, P. Liang, M. Shahin, A. Di Salle, and G. Márquez, "Design, monitoring, and testing of microservices systems: The practitioners' perspective," *Journal of Systems and Software*, vol. 182, p. 111061, 2021.

[11] J. Bogner, J. Fritzsch, S. Wagner, and A. Zimmermann, "Microservices in industry: insights into technologies, characteristics, and software quality," in *Proceedings of the 16th IEEE international conference on software architecture companion (ICSA-C)*. IEEE, 2019, pp. 187–195.

[12] A. Rezaei Nasab, M. Shahin, P. Liang, M. E. Basiri, S. A. H. Raviz, H. Khalajzadeh, M. Waseem, and A. Naseri, "Automated identification of security discussions in microservices systems: Industrial surveys and experiments," *Journal of Systems and Software*, vol. 181, p. 111046, 2021.

[13] A. Pereira-Vale, E. B. Fernandez, R. Monge, H. Astudillo, and G. Márquez, "Security in microservice-based systems: A multivocal literature review," *Computers & Security*, vol. 103, p. 102200, 2021.

[14] M. Waseem, P. Liang, M. Shahin, A. Ahmad, and A. Rezaei Nasab, "On the nature of issues in five open source microservices systems: An empirical study," in *Proceedings of the 25th International Conference on Evaluation and Assessment in Software Engineering (EASE)*. ACM, 2021, pp. 201–210.

[15] J. Soldani, D. A. Tamburri, and W.-J. Van Den Heuvel, "The pains and gains of microservices: A systematic grey literature review," *Journal of Systems and Software*, vol. 146, pp. 215–232, 2018.

[16] T. Yarygina and A. H. Bagge, "Overcoming security challenges in microservice architectures," in *Proceedings of the 12th IEEE Symposium on Service-Oriented System Engineering (SOSE)*. IEEE, 2018, pp. 11–20.

[17] A. Hannousse and S. Yahiouche, "Securing microservices and microservice architectures: A systematic mapping study," *Computer Science Review*, vol. 41, p. 100415, 2021.

[18] N. C. Mendonça, C. Box, C. Manolache, and L. Ryan, "The monolith strikes back: Why istio migrated from microservices to a monolithic architecture," *IEEE Software*, vol. 38, no. 5, pp. 17–22, 2021.

[19] V. Lenarduzzi, F. Lomio, N. Saarimäki, and D. Taibi, "Does migrating a monolithic system to microservices decrease the technical debt?" *Journal of Systems and Software*, vol. 169, p. 110710, 2020.

[20] M.-O. Pahl, F.-X. Aubet, and S. Liebald, "Graph-based IoT microservice security," in *Proceedings of the IEEE/IFIP Network Operations and Management Symposium (NOMS)*. IEEE, 2018, pp. 1–3.

[21] A. Rezaei Nasab, M. Shahin, S. A. Hoseyni Raviz, P. Liang, A. Mashmool, and V. Lenarduzzi. (2021) dataset. [Online]. Available: https://doi.org/10.5281/zenodo.5791337

[22] C. Pahl, P. Jamshidi, and O. Zimmermann, "Architectural principles for cloud software," *ACM Trans. Internet Technol.*, vol. 18, no. 2, p. Article No.: 17, 2018.

[23] D. Taibi, V. Lenarduzzi, and C. Pahl, "Processes, motivations, and issues for migrating to microservices architectures: An empirical investigation," *IEEE Cloud Computing*, vol. 4, no. 5, pp. 22–32, 2017.

[24] J. Soldani, D. A. Tamburri, and W.-J. V. D. Heuvel, "The pains and gains of microservices: A systematic grey literature review," *Journal of Systems and Software*, vol. 146, pp. 215–232, 2018.

[25] J. A. Scott, *A Practical Guide to Microservices and Containers*. MapR Data Technologies, 2018.

[26] K. Torkura, M. I. H. Sukmana, and C. Meinel, "Integrating continuous security assessments in microservices and cloud native applications," in *Proceedings of the 10th International Conference on Utility and Cloud Computing (UCC)*. ACM, 2017, pp. 171–180.

[27] F. Ponce, J. Soldani, H. Astudillo, and A. Brogi, "Smells and refactorings for microservices security: A multivocal literature review," *arXiv preprint arXiv:2104.13303*, 2021.

[28] M.-O. Pahl and L. Donini, "Securing IoT microservices with certificates," in *Proceedings of the IEEE/IFIP Network Operations and Management Symposium (NOMS)*. IEEE, 2018, pp. 1–5.

[29] D. Yu, Y. Jin, Y. Zhang, and X. Zheng, "A survey on security issues in services communication of microservices-enabled fog applications," *Concurrency and Computation: Practice and Experience*, vol. 31, no. 22, p. e4436, 2019.

[30] N. Chondamrongkul, J. Sun, and I. Warren, "Automated security analysis for microservice architecture," in *Proceedings of the 17th IEEE International Conference on Software Architecture Companion (ICSA-C)*. IEEE, 2020, pp. 79–82.

[31] Y. Sun, S. Nanda, and T. Jaeger, "Security-as-a-service for microservices-based cloud applications," in *Proceedings of the 7th IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*. IEEE, 2015, pp. 50–57.

[32] T. Bi, P. Liang, A. Tang, and X. Xia, "Mining architecture tactics and quality attributes knowledge in stack overflow," *Journal of Systems and Software*, vol. 180, p. 111005, 2021.

[33] I. Malavolta, G. A. Lewis, B. Schmerl, P. Lago, and D. Garlan, "Mining guidelines for architecting robotics software," *Journal of Systems and Software*, vol. 178, p. 110969, 2021.

[34] N. Meng, S. Nagy, D. Yao, W. Zhuang, and G. A. Argoty, "Secure coding practices in java: Challenges and vulnerabilities," in *Proceedings of the 40th International Conference on Software Engineering (ICSE)*. ACM, 2018, pp. 372–383.

[35] B. G. Glaser, A. L. Strauss, and E. Strutzel, "The discovery of grounded theory; strategies for qualitative research," *Nursing Research*, vol. 17, no. 4, p. 364, 1968.

[36] B. A. Kitchenham and S. L. Pfleeger, "Personal opinion surveys," in *Guide to Advanced Empirical Software Engineering*. Springer, 2008, pp. 63–92.

[37] (2021, July) Single sign-on in microservice architecture. [Online]. Available: https://stackoverflow.com/questions/25595492
[38] (2021, July) Authorization between services. [Online]. Available: https://github.com/moleculerjs/moleculer/issues/304
[39] (2021, July) Identity/customer service as a microservice. [Online]. Available: https://github.com/dotnet-architecture/eShopOnContainers/issues/785
[40] (2021, July) Single sign on: Azure ad b2c vs identityserver4, and others. [Online]. Available: https://github.com/dotnet-architecture/eShopOnContainers/issues/925
[41] (2021, July) How to refresh request token with microservice multiple instances? [Online]. Available: https://stackoverflow.com/questions/48760583
[42] (2021, July) Discussion on security. [Online]. Available: https://github.com/spinnaker/spinnaker/issues/148
[43] (2018, June) Should api gateway be responsible for authorisation? [Online]. Available: https://stackoverflow.com/questions/50700178
[44] (2017, June) Securing ui of jaeger. [Online]. Available: https://github.com/jaegertracing/jaeger/issues/218
[45] (2021, July) Micro-service architecture, should the spring cloud config server, zuul gateway server and eureka server be protected as resources? [Online]. Available: https://stackoverflow.com/questions/61307668
[46] (2021, July) Proposal: Create the template function for authentication in the file for each service. [Online]. Available: https://github.com/goadesign/goa/issues/2361
[47] (2021, July) How to authenticate json web tokens (jwt) across different apis? [Online]. Available: https://stackoverflow.com/questions/55394912
[48] (2021, July) gateway api. [Online]. Available: https://github.com/dotnet-architecture/eShopOnContainers/issues/239
[49] (2021, July) Laravel passport, oauth and microservices. [Online]. Available: https://stackoverflow.com/questions/39126827
[50] (2021, July) Should jwt be a separate auth micro-service and not sit with the backend business logic? [Online]. Available: https://stackoverflow.com/questions/58948497
[51] (2019, August) Rfc: Allow spring property placeholders in pipeline expressions. [Online]. Available: https://github.com/spinnaker/spinnaker/issues/4725
[52] (2021, July) Startup.cs - add authorization with ocelot. [Online]. Available: https://github.com/dotnet-architecture/eShopOnContainers/issues/647
[53] (2021, July) In a microservice environment, should any producer be able to verify jwt tokens? [Online]. Available: https://stackoverflow.com/questions/44471051
[54] (2021, July) Shared signature key for jwt in various microservices. [Online]. Available: https://stackoverflow.com/questions/43559197
[55] (2018, November) Rfc: Halyard secret management. [Online]. Available: https://github.com/spinnaker/spinnaker/issues/3649
[56] (2020, February) Hide passwords in urls on the /config endpoint. [Online]. Available: https://github.com/cortexproject/cortex/pull/2176
[57] (2019, March) Vulnerable data exposed with metrics endpoint. [Online]. Available: https://github.com/jaegertracing/jaeger/issues/1428
[58] (2019, September) grpc plugin framework does not respect –query.bearer-token-propagation flag. [Online]. Available: https://github.com/jaegertracing/jaeger/issues/1821
[59] (2021, January) Cortex feature request/improvement - refresh aws object store credentials for expired tokens. [Online]. Available: https://github.com/cortexproject/cortex/issues/3731
[60] (2021, July) Microservices - how to solve security and user authentication? [Online]. Available: https://stackoverflow.com/questions/32574103
[61] (2021, July) Decoding oauth2 jwt at api gateway level vs at individual microservice level. [Online]. Available: https://stackoverflow.com/questions/51524648
[62] (2021, July) Find the best location to inject server information to the routing handler. [Online]. Available: https://github.com/networknt/light-4j/issues/11
[63] (2021, July) Add logging module for light 4j rfc#29. [Online]. Available: https://github.com/networknt/light-4j/issues/453
[64] Introduction to json web tokens. [Online]. Available: https://jwt.io/introduction
[65] (2021, July) How is https/ssl termination handled? [Online]. Available: https://github.com/dotnet-architecture/eShopOnContainers/issues/752
[66] (2021, July) If we have already implemented the authorization in .net core micro-service api gateway do we need to implement in all micro services as well? [Online]. Available: https://stackoverflow.com/questions/54631411
[67] (2020, August) Iam authentication support in ruler and alertmanager s3 client. [Online]. Available: https://github.com/cortexproject/cortex/issues/3034
[68] (2020, October) Authenticating to gcp when using chunks storage (bigtable and gcs). [Online]. Available: https://github.com/cortexproject/cortex/issues/3306
[69] (2021, July) Oauth2 grant for server-to-server communication. [Online]. Available: https://stackoverflow.com/questions/27838280
[70] (2021, July) Private services. [Online]. Available: https://github.com/moleculerjs/moleculer/issues/124
[71] (2021, July) How to authenticate and authorize in a microservice architecture? [Online]. Available: https://stackoverflow.com/questions/52349986
[72] (2021, July) Addtocart method relies on the posted productdetails. [Online]. Available: https://github.com/dotnet-architecture/eShopOnContainers/issues/250
[73] (2021, July) How to add an api with oauth2 on the top of kong? [Online]. Available: https://stackoverflow.com/questions/47324184
[74] (2021, July) Quick solution to handle service to service authentication in a microservices architecture. [Online]. Available: https://stackoverflow.com/questions/61433192
[75] (2017, September) Span authentication support in jaeger collector. [Online]. Available: https://github.com/jaegertracing/jaeger/issues/427
[76] (2020, November) Flaky test: Testreload. [Online]. Available: https://github.com/jaegertracing/jaeger/issues/2622
[77] (2021, July) Microservices and database security. [Online]. Available: https://stackoverflow.com/questions/53621693
[78] (2020, August) Add tls client reload. [Online]. Available: https://github.com/cortexproject/cortex/issues/3012
[79] (2021, July) Oauth 2.0 in microservices: When a resource server communicates with another resource server. [Online]. Available: https://stackoverflow.com/questions/52290697
[80] (2021, July) eshoponcontainers. [Online]. Available: https://github.com/dotnet-architecture/eShopOnContainers/wiki/gRPC
[81] (2021, July) Jaeger. [Online]. Available: https://www.jaegertracing.io/docs/1.27/architecture/
[82] (2018, July) Jaeger trace sampling should not be decided by every service (by default). [Online]. Available: https://github.com/cortexproject/cortex/issues/885
[83] (2017, October) Build a secure channel for security reports. [Online]. Available: https://github.com/jaegertracing/jaeger/issues/457
[84] (2017, October) Allow secure communication between components. [Online]. Available: https://github.com/jaegertracing/jaeger/issues/458
[85] (2021) Azure data security and encryption best practices. [Online]. Available: https://docs.microsoft.com/en-us/azure/security/fundamentals/data-encryption-best-practices#protect-data-in-transit
[86] I. Grigorik and Surma. Http/2. [Online]. Available: https://developers.google.com/web/fundamentals/performance/http2
[87] Google. Protocol buffers. [Online]. Available: https://developers.google.com/protocol-buffers
[88] (2018, May) Deploying spinnaker with halyard to k8s with kube v2 provider and ssl enabled for gate fails because k8s readinessprobe fails. [Online]. Available: https://github.com/spinnaker/spinnaker/issues/2765
[89] (2021, July) Login authentication flow for microservices. [Online]. Available: https://stackoverflow.com/questions/59058573
[90] (2021, July) After customerbasket has been posted to basketcontroller where is the unitprice validated with the catalog in the workflow? [Online]. Available: https://github.com/dotnet-architecture/eShopOnContainers/issues/945
[91] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén, *Experimentation in Software Engineering*. Springer, 2012.