

# An Empirical Study of Security Practices for Microservices Systems

Rezaei Nasab Ali<sup>a</sup>, Shahin Mojtaba<sup>b</sup>, Hoseyni Raviz Seyed Ali<sup>a</sup>, Liang Peng<sup>a,\*</sup>, Mashmool Amir<sup>c</sup> and Lenarduzzi Valentina<sup>d</sup>

<sup>a</sup>School of Computer Science, Wuhan University, 430072 Wuhan, China

<sup>b</sup>School of Computing Technologies, RMIT University, 3000 Melbourne, Australia

<sup>c</sup>Department of Computer Science, Bioengineering, Robotics and System Engineering, University of Genoa, 16126 Genoa, Italy

<sup>d</sup>Faculty of Information Technology and Electrical Engineering, University of Oulu, FI-90014 Oulu, Finland

## ARTICLE INFO

### Keywords:

Microservice  
Security  
Empirical Study  
Practitioners  
Practice

## 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. We first manually analyzed 861 microservices security points, including 567 issues, 9 documents, and 3 wiki pages from 10 GitHub open-source microservices systems and 306 Stack Overflow posts concerning security in microservices systems. In this study, a microservices security point is referred to as “a GitHub issue, a Stack Overflow post, a document, or a wiki page that entails 5 or more microservices security paragraphs”. Our analysis led to a catalog of 28 microservices security practices. We then ran a survey with 74 microservices practitioners to evaluate the usefulness of these 28 practices. Our findings demonstrate that the survey respondents affirmed the usefulness of the 28 practices. 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.

## 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-tolerance [1, 3].

Since the advent of the MSA style, securing microservices systems has been a challenge for software practitioners and organizations [4, 5, 6, 7, 8, 9, 10]. The potential security challenges associated with microservices systems may compel software organizations to revisit their decision to adopt or migrate to microservices [11, 12]. 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

[13, 14, 15, 16]; (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., [6, 9, 17]), and recent review studies have called for more studies on security in microservices systems [1, 18, 8].

While these works are valuable for security in microservices systems, documented knowledge and guidelines on how software practitioners design and implement secure microservices systems are scarce (if any) [6]. It is argued that software practitioners are keen to learn and apply the practices and decisions that their peers used or are using during their development process [19, 20, 14]. Software practitioners are also interested in learning bad practices adopted by their peers to avoid repeating mistakes. Hezavei et al. [20] defined a software development practice as “an activity or step carried out to achieve a goal during the development process”. To the best of our knowledge, no systematic attempt has been made to develop and document a catalogue of best practices used by software practitioners that enable the development of secure microservices systems. Hence, the ultimate goal of this study is to empirically collect and document microservices-specific security practices. At the same time, it is also important to understand how practitioners perceive the usefulness of these security best practices. To this end, we define the following two research questions (RQs).

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

**RQ2. To what extent do practitioners consider the iden-**

\*Corresponding author

✉ rezaei.ali.nasab@gmail.com (R.N. Ali);

mojtaba.shahin@rmit.edu.au (S. Mojtaba); s.ali.hoseyni@gmail.com (H.R. Seyed Ali); liangp@whu.edu.cn (L. Peng); 5245307@studenti.unige.it (M. Amir); valentina.lenarduzzi@oulu.fi (L. Valentina)

ORCID(s):

### *tified microservices security practices useful?*

To answer these two RQs, we conducted an empirical study that identified and validated 28 security practices for microservices systems. We first collected and manually analyzed 861 microservices security points, including 567 issues, 9 documents, and 3 wiki pages from 10 GitHub open-source microservices systems and 306 Stack Overflow posts concerning security in microservices systems (see Section 3.1). A microservices security point is defined as “a GitHub issue, a Stack Overflow post, a document, or a wiki page that entails 5 or more microservices security paragraphs”. This manual analysis led to a catalog of 28 microservices security practices. The security practices identified from GitHub and Stack Overflow are based on the authors’ analysis, which might be subjective and unreliable. Consequently, we ran an industrial survey completed by 74 practitioners to seek their perceptions about the usefulness of the identified security practices (see Section 3.2).

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

- Identification of 28 security practices in 6 categories for microservices systems;
- Validation of the usefulness of these security practices from 74 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].

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, 18]. For example, Balalaie et al. [23] identified 15 patterns for this purpose, and Auer et al. [24] developed an assessment framework to help software organizations specify and measure possible advantages and difficulties of migration to microservices. However, the migration to or adoption of microservices can be associated with many serious challenges and issues, which need careful consideration [8, 25]. 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 [11, 12].

Some researchers also looked at approaches and tools for microservices systems. Cinque et al. [26] developed an approach for monitoring microservices, and Heorhiadi et al. [27] introduced a framework, *Gremlin*, to test the “failure-handling capabilities” of microservices. Other researchers empirically investigated how microservices systems are designed, implemented, tested, and monitored in the software industry (e.g., [28, 29]).

#### **2.1.2. Security in Microservices Systems**

Besides the various advantages brought by microservices, security often becomes an issue during their deployment. Similar to other types of systems (e.g., monolithic systems), improved security in microservices systems may be accomplished in different ways, such as by applying a secure development methodology [30, 31, 32]. 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 [7]. 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<sup>1</sup> or Istio<sup>2</sup> 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 ser-

<sup>1</sup><https://kubernetes.io>

<sup>2</sup><https://istio.io>

vices [7].

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 [7]. Moreover, developers need to discriminate how to update the authentication mechanisms every time when new services or new users are included in the system.

The security and authentication issues 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 [33]. An attempt to reduce this issue is provided by the “Moving Target Defenses”, which proposes to modify component images to create uncertainty for attackers [34].

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 [8, 25]. This requires creating a secure communication channel between the monolithic system and the microservices system and integrating the security and authentication system adopted for the monolithic system with the new one adopted in the microservices system [8, 25]. 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

Researchers and practitioners may use security techniques, tools, practices, patterns, or secure development methodologies to support the development of secure software systems [30, 31, 32]. This section covers the studies that have targeted security in microservices systems.

Several secondary studies have been conducted on microservices systems and on security in microservices. In a recent review study, Waseem et al. [18] observed that there are a few concrete solutions for addressing security concerns when implementing microservices systems in DevOps. Another systematic mapping review on 46 papers by Hannousse and Yahiouche [10] 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. [7] 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. [35] focused on security smells and refactorings and collected and analyzed 58 white and grey literature published from 2014 to 2021. They found 10 security smells with their security properties and corresponding refactorings. They also elaborated on how the corresponding refactorings mitigate the effects of the smells. A recent systematic grey literature study by Billawa et al. on 57 microservices-related grey literature sources identified 7 security challenges, including “*trust between services*”, “*large attack area*”, “*testing*”, “*container management*”, “*low visibility*”, “*secret management*”, and “*polyglot architecture*” [36]. Billawa et al. also found that the practices, such as “*defense in depth*”, “*DevSecOps*”, “*encrypt sensitive data*”, “*immutable container*”, “*rate throttling*”, “*secure-by-design*”, and “*least privilege*” can help develop secure microservices systems.

Waseem et al. [5] empirically investigated the issues reported in 5 open source microservices systems hosted on GitHub. They found that 10.18% of these issues can be attributed to security. Yarygina and Bagge [9] 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 6 layers (hardware, virtualization, cloud, communication, service/application, and orchestration) with solutions to address them (e.g., secure implementation of service discovery and registry components). Richardson [37] introduced 44 microservices patterns that are categorized into 16 groups (e.g., “data consistency” patterns). Among them, “Access Token” is the only pattern related to security.

To secure IoT microservices, Pahl et al. [17] 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 [38] 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. [39] 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 4 groups: containers issues, data issues, permission issues, and network issues.

By surveying 67 participants, Rezaei Nasab et al. [6] affirmed that security challenges in microservices systems differ from non-microservices systems. To bridge the security knowledge gap among microservices practitioners [13, 14, 15, 16], they developed a set of machine and deep learn-

ing approaches to automatically recognize security discussions (including security design decisions, challenges, or solutions) from open source microservices systems. Chondamrongkul et al. [40] 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. [41] 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 trusted and ready-to-use security best practices for microservices systems. For example, although practices identified in [36] are valuable, they are still general and need concrete solutions to be implemented in microservices system development. Our 28 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 74 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. In this section, we first explain our approach to identifying 28 security practices for microservice systems from GitHub and Stack Overflow (RQ1) in Section 3.1. We then discuss the design and execution of a survey to evaluate these security practices (RQ2) in Section 3.2.

#### 3.1. Mining Security Practices (RQ1)

In this section, we first describe how we collected data related to microservices security from GitHub and Stack Overflow (see Section 3.1.1), and then how the collected data was analyzed to identify microservices security practices (see Section 3.1.2).

##### 3.1.1. Data Collection

GitHub issues and Stack Overflow posts (including Stack Exchange) are valuable sources for identifying development practices (e.g., architectural practices [42, 43] and security practices [44]). Hence, we focused on GitHub and Stack Overflow to collect data related to microservices security.

**Step 1.** In our previous work [6], we manually created a dataset of 5,018 security paragraphs collected from 567 GitHub issues in 7 open-source microservices systems and 505 Stack Overflow posts. These security paragraphs include “*design decisions, challenges, or solutions relating to security in microservices systems*” [6]. The architectural style of these 7 open-source systems (i.e., goa, eShopOnContainers, microservices-demo, scalecube-services, molecular, 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 3 contributors who have the most commits in a project*” [6]. Our previous work [6] aimed to develop ML/DL-based approaches to automatically differentiate security paragraphs from non-security paragraphs in GitHub issues and Stack Overflow posts concerning security in microservices. In contrast, this work aims to identify security best practices for microservices systems from GitHub and Stack Overflow and then evaluate them with software practitioners.

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 GitHub issues or Stack Overflow posts that entail equal or more than 5 such security paragraphs. Our decision to set the threshold to 5 security paragraphs was based on our experience in the previous work [6] as issues and posts with 5 or more security paragraphs are more suitable for identifying security practices. As shown in Figure 1, this process led to 76 candidate GitHub issues and 306 candidate Stack Overflow posts, which are called security points. The creation date of the 306 candidate Stack Overflow posts ranges from August 2014 to July 2020.

**Step 2.** Apart from the 7 open-source microservices systems used in our previous work [6], we found 3 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 3 new systems also employ the MSA style. We used the same approach in [6] (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 [6]) to identify security paragraphs from the GitHub issues of these 3 new projects. Since the unit of analysis in the DeepM1 approach [6] was a paragraph, we converted the GitHub issues of these 3 projects to paragraphs using their HTML tag <p>. The pre-processing used in [6] 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 the GitHub issues of these 3 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 only selected the GitHub issues of these 3 projects that entail equal or more than 5 such security paragraphs. This resulted in a collection of 467 security points from spinnaker, cortex, and jaeger projects (see Figure 1).

**Step 3.** Among all systems discussed in **Step 1** and **Step 2**, 9 systems (i.e., eShopOnContainers, spinnaker, cortex, jaeger, goa, molecular, deep-framework, microservices-demo, and light-4j) have documentation with 5 or more security paragraphs. Further, 3 systems, including eShopOnContain-






Figure 1: An overview of extracting microservices security practices

ers, scalecube-services, and light-4j, have a wiki page with 5 or more security paragraphs. We considered these 9 pieces of documentation and 3 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 3 phases.

**Pilot Phase.** We randomly selected 20 security points from the 861 security points and asked 3 analysts (3 authors) to analyze them independently. The goal was to get familiar with data and understand what sort of security points should be considered a security practice. Each analyst applied the open coding and constant comparison techniques from Grounded Theory [45] 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 3 analysts collaboratively analyzed the remaining 841 security points. 90 security points were allo-

cated to the 3 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 [45]. 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, which were grouped into 6 categories based on their topics: *Authorization and Authentication, Token and Credentials, Internal and External Microservices, Microservices Communications, Private Microservices, and Database and Environments*.

**Feedback Phase.** In this phase, the other 3 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 4 practices with other microservices security 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 com-

**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 <= 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                 |
| "How many years of experience do you have with security in microservices systems?" | Multiple choice | 0 <year <= 1                 |
| "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                    |

pleted by 6 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 2 security practices did not contain enough information to be understood. We also found 2 practices that overlapped and merged them. We further slightly rephrased the wording of some practices (e.g., adding more information to a security practice) to remove any ambiguities. These 28 security practices are distributed into the 6 categories developed in the **Main Phase** as follows: *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 a survey (called the validation survey) to evaluate the usefulness of the 28 security practices for microservices systems collected in Section 3.1.

#### 3.2.1. Protocol

Considering the guidelines proposed by Kitchenham and Pfleege [46], 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 3 groups.

**Demographic questions.** We asked 7 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 7 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 prac-

tices using a Likert question. The mandatory Likert scale questions were ranked on a 4-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 6 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, 2 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 10 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 74 valid responses. The initial analysis of the survey responses revealed that 2 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).

**Table 2**

A list of 10 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  | <a href="https://bit.ly/2X40b4M">https://bit.ly/2X40b4M</a> | 18.6k | 7.9k  | 1,752  | 17   | 142   | 120k | C#         | ✓    | ✓    |
| 2  | jaeger             | <a href="https://bit.ly/2YvyJgN">https://bit.ly/2YvyJgN</a> | 14.2k | 1.7k  | 3,222  | 44   | 221   | 105k | Go, Shell  | ✓    | ✗    |
| 3  | spinnaker          | <a href="https://bit.ly/3ndjJyu">https://bit.ly/3ndjJyu</a> | 8k    | 1.1k  | 6,514  | 132  | 116   | 6k   | Shell, Go  | ✓    | ✗    |
| 4  | moleculer          | <a href="https://bit.ly/2X5DayD">https://bit.ly/2X5DayD</a> | 4.6k  | 441   | 1,003  | 99   | 92    | 98k  | Javascript | ✓    | ✗    |
| 5  | goa                | <a href="https://bit.ly/38LUYkD">https://bit.ly/38LUYkD</a> | 4.4k  | 459   | 2,907  | 55   | 88    | 88k  | Go         | ✓    | ✗    |
| 6  | cortex             | <a href="https://bit.ly/3nbq65x">https://bit.ly/3nbq65x</a> | 4.3k  | 604   | 4,438  | 50   | 204   | 1.2m | Go         | ✓    | ✗    |
| 7  | light-4j           | <a href="https://bit.ly/3zOnhL0">https://bit.ly/3zOnhL0</a> | 3.3k  | 554   | 1,030  | 140  | 34    | 57k  | Java       | ✓    | ✓    |
| 8  | microservices-demo | <a href="https://bit.ly/38LtxY2">https://bit.ly/38LtxY2</a> | 2.8k  | 1.8k  | 875    | 13   | 55    | 15k  | Python     | ✓    | ✗    |
| 9  | deep-framework     | <a href="https://bit.ly/3ncbgM4">https://bit.ly/3ncbgM4</a> | 538   | 75    | 647    | 22   | 12    | 956k | Javascript | ✓    | ✗    |
| 10 | scalecube-services | <a href="https://bit.ly/3DTFhGn">https://bit.ly/3DTFhGn</a> | 507   | 79    | 820    | 178  | 21    | 11k  | Java       | ✗    | ✓    |



**Figure 2:** Experience of the participants ( $n=74$ ) 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), and organization size of the participants (E).

### 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 [45]. 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


We provide the demographics of the survey respondents.

**Experience.** Figure 2 (A) shows that 51.4% of the participants ( $n=74$ ) have been involved in software development for at least 10 years. All participants, except for one partic-


ipant, had at least one year of experience in developing microservices systems, with 50.1% having more than 3 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. 25.7% worked with security in microservices systems in less than one year. The rest (4.1%) did not have any experience in this regard.

**Role.** As shown in Figure 2 (D), the participants mainly worked as Developer (28.4%, 21 out of 74), Architect (24.3%, 18 out of 74), Software Engineer (21.6%, 16 out of 74), and Technical Lead (14.9%, 11 out of 74).

**Organization size and domain.** The majority of the participants (83.8%) came from organizations with more than 100 employees (see Figure 2 (E)). 59.5% (44 out of 74 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



**Figure 3:** Participants' organization domains (n=74). Note: Participants could select more than one domain




**Figure 4:** Number of participants (n=74) from 23 countries

organization domains in the demographic question. The dominant domains are “consulting and IT services” and “financial”, 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 50 responses for this question. The 50 participants who indicated their country information came from 23 countries across 6 continents, including Europe (10 countries), Asia (6 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).

## 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 6 groups. Similar to the arguments by Malavolta et al. [43], 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



**Figure 5:** The level of usefulness of the 28 identified security practices for microservices systems from the perspective of 74 practitioners

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

### 4.2.1. Authorization and Authentication

This group includes 6 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) [7]. PA1 is a way to implement SSO [47]. 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. 58 (78.38%) out of the 74 survey respondents



**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

| Security practices for authorizing and authentication |   | Sources  | AU    | U     | NU    | ANU   | IDK   | MED | AVG  |
|---|---|--|-------|-------|-------|-------|-------|-----|------|
| <b>PA1</b>  | 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.  | [47], [48], [49], [50], [51], [52], [53], [54] | 32.43 | 45.95 | 16.22 | 1.35  | 4.05  | 3   | 3.14 |
| <b>PA2</b>  | Each microservice in the microservices architecture must be responsible for its own security, i.e., each microservice must have security enabled.   | [47], [55], [56], [57]                         | 40.54 | 33.78 | 17.57 | 6.76  | 1.36  | 3   | 3.10 |
| <b>PA3</b>  | 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 <b>PA2</b> ) because the internal microservices can be protected by not being published out of the Docker Host. | [58]   | 28.39 | 24.32 | 24.32 | 12.16 | 10.81 | 3   | 2.77 |
| <b>PA4</b>  | A large microservices system is recommended to use an API Gateway approach for securing/authorizing/routing microservices.  | [58], [57], [59], [60], [61]                   | 70.27 | 25.68 | 4.05  | 0.00  | 0.00  | 4   | 3.66 |
| <b>PA5</b>  | 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.   | [62]   | 45.95 | 40.54 | 12.16 | 0.00  | 1.35  | 3   | 3.34 |
| <b>PA6</b>  | Use Public Key Infrastructure (PKI) signing/verification system to prevent round trips to the Authorization Service.  | [63]   | 21.63 | 47.30 | 14.86 | 1.35  | 14.86 | 3   | 3.05 |

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 [47]. In other words, with this practice, the user will have to authorize each microservice individually. More than half of the respondents (74.32%) acknowledged that the practice is *absolutely useful* or *useful*. The follow-

ing are 4 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)

👍 "Absolutely useful **as long as microservices don't bother each other with different security protocols**. In my company we used to secure every module separately to ensure powerful security while maintaining their best performance." (Developer)

👍 "Enabling security between microservices - **this depends on the use case**. But it is **additional overhead**. If we are isolated from APIGW and the microservices needs to communicate with any microservices and if protected, then only it is advisable to have the security between microservices." (Architect)

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 [58]. 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. 39 survey participants (52.71%) considered this practice (*absolutely*) *useful*. On the other hand, 27 (36.48%) opposed it (24.32% rated it as *not useful*, and 12.16% rated it as *absolutely not useful*). The following is a positive comment that we received from the respondents.

👍 “If we make the **security check at api gateway level**, then the microservices behind the gateway can **ignore token validity check** request to auth server. This will **reduce overhead**. Though they need to **check the access level for maintaining Role-Based Access Control (RBAC)**.” (Technical Lead)

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)

👍 “It is important to **secure communication between microservices** even if you use API Gateway.” (Technical Lead)

💡 **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 [57, 58]. Out of 28 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.

👍 “[This practice] **simplify code and is robust**”

👍 “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<sup>3</sup> [62]. However, developers need to balance between security and simplicity. The majority of our survey respondents agreed with the usefulness of this practice (86.49% 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 [63]. Most respondents (68.93% 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.

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

**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

| Security practices for tokens and credentials |  | Sources                      | AU    | U     | NU    | ANU   | IDK   | MED | AVG  |
|---|--|------------------------------|-------|-------|-------|-------|-------|-----|------|
| <b>PT1</b>                                    | Use a method based on the Public/Private key to secure microservices through JSON Web Token (JWT).                                 | [64], [65], [66], [67], [68] | 39.19 | 51.35 | 2.70  | 0.00  | 6.76  | 3   | 3.39 |
| <b>PT2</b>                                    | In microservices systems, use JSON Web Tokens (JWTs) to handle the session expiration/revocation.                                  | [47], [69]                   | 40.54 | 48.65 | 4.05  | 0.00  | 6.76  | 3   | 3.39 |
| <b>PT3</b>                                    | Secure caches of credentials in each microservice that needs to access other microservices.  | [70]                         | 14.87 | 43.24 | 20.27 | 10.81 | 10.81 | 3   | 2.70 |
| <b>PT4</b>                                    | Decode JSON Web Token (JWT) at the microservices level instead of the API Gateway level.   | [71]                         | 27.03 | 33.78 | 18.92 | 6.76  | 13.51 | 3   | 2.94 |
| <b>PT5</b>                                    | Endpoints of microservices like server information, health check, and logging level must be secured in the request/response chain. | [72], [73], [62]             | 36.49 | 35.14 | 18.91 | 2.70  | 6.76  | 3   | 3.13 |

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

☞ “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 5 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 2 key pairs (private signing key and public verification key) [74]. 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 [64]. If the microservices system uses the API Gateway approach, the API Gateway should also know the private signing key.

90.54% (67) of the survey respondents rated it as *absolutely useful* or *useful*. Only 2 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<sup>4</sup> to track token revocations [47]. **PT2** received almost similar positive feedback to **PT1** from the survey respondents (89.19% *absolutely useful* or *useful*). A respondent stated that this practice is use-

ful 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** [70] is *absolutely useful* or *useful*, while 31.08% rated it as *not useful* or *absolutely not useful*. A software engineer with more than 3 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 whether 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. 45 participants

<sup>4</sup><https://redis.io>

(60.81%) 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 [71]. Our analysis shows that 19 respondents (25.68%) did not agree with the usefulness of **PT4**. Below are 3 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)

🗨️ “I think that JWT tokens can be checked at the API Gateway, **but not 'instead of' the microservice level**. Checking at the gateway will keep failed authentications away from any unnecessary processing but they should be verified by the Microservice.” (Technical Lead)

💡 **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 [72, 73]. Depending on the requirements of a microservices system, developers may only use some of these endpoints. 71.63% 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).

💡 **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. 52 respondents (70.27%) acknowledged that developers could use unique tokens for securing internal and external microservices. Two practitioners rated it *absolutely not useful* (2.7%), and 12 practitioners considered it *not useful* (16.22%). Below are 2 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 [75]. 60.81% of the participants stated that **PI2** is *absolutely useful* or *useful*. 20.27% 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)

🗨️ “External authorization and Internal authorization are different. **External token must be used to authorize the user. Don’t mix**.” (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 (83.78%, 62 out of 74) 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 [76].

👍 “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.



**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

| Security practices for internal and external microservices |   | Sources          | AU    | U     | NU    | ANU  | IDK   | MED | AVG  |
|--|---|------------------|-------|-------|-------|------|-------|-----|------|
| <b>PI1</b>   | 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.  | [75]             | 27.03 | 43.24 | 16.22 | 2.7  | 10.81 | 3   | 3.06 |
| <b>PI2</b>   | 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. | [75]             | 17.57 | 43.24 | 13.51 | 6.76 | 18.92 | 3   | 2.88 |
| <b>PI3</b>   | Whether microservices are only internally used within an organization or are externally accessible to third parties, authentication is required either way.   | [76], [77], [78] | 54.05 | 29.73 | 6.76  | 5.41 | 4.05  | 4   | 3.38 |
| <b>PI4</b>   | In an internal microservice use case, "client credential" should not get exposed to the third party.  | [79]             | 56.76 | 35.14 | 4.05  | 1.35 | 2.70  | 4   | 3.51 |
| <b>PI5</b>   | Microservices systems made of components should be isolated and internal calls should not be leaked outside their boundaries.   | [80]             | 51.35 | 40.54 | 5.41  | 0.00 | 2.70  | 4   | 3.47 |
| <b>PI6</b>   | Encrypt tokens if they are going to be exposed to the outside of the system boundary.   | [81]             | 55.41 | 22.97 | 12.16 | 4.05 | 5.41  | 4   | 3.37 |
| <b>PI7</b>   | 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.  | [82]             | 37.84 | 39.19 | 13.51 | 0.00 | 9.46  | 3   | 3.27 |

🗨️ **"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 [79]. 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 (91.89%, 68 out of 74) 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 [81]. This practice was *absolutely useful* for 41 respondents (55.41%) and *useful* for 17 respondents (22.97%). The following are 2 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 [82]. It also impacts the performance

**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

| Security practices for microservices communications |  | Sources                            | AU    | U     | NU    | ANU  | IDK   | MED | AVG  |
|---|--|------------------------------------|-------|-------|-------|------|-------|-----|------|
| <b>PM1</b>  | Use OAuth2 "Client Credentials Flow" if two microservices that trust each other want to talk together from the backends.   | [83], [84], [85], [86]             | 24.32 | 50.0  | 6.76  | 2.70 | 16.22 | 3   | 3.15 |
| <b>PM2</b>  | The connection between a microservice and its respective database should be protected by a security protocol, like Transport Layer Security (TLS).               | [87], [88]                         | 48.65 | 36.49 | 8.11  | 2.70 | 4.05  | 4   | 3.37 |
| <b>PM3</b>  | When a microservice in a microservices architecture needs to call another microservice, the access token should be passed around microservices with the request. | [89]                               | 35.14 | 37.84 | 18.92 | 4.05 | 4.05  | 3   | 3.08 |
| <b>PM4</b>  | It is recommended to use the gRPC framework for internal microservice-to-microservice synchronous communication.   | [86], [90], [91], [92], [93], [94] | 20.27 | 37.84 | 16.22 | 0.00 | 25.67 | 3   | 3.05 |

of microservices when one of them does not perform well [82].

57 (77.03%) 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<sup>5</sup>.

👍 "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)

👍 "This is required for **performance**."

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 4 practices that are related to authenticating and authorizing requests when 2 or more microservices are communicating.

<sup>5</sup><https://samnewman.io/patterns/architectural/bff>

💡 **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 [83]. Assume that there are a user and 2 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 2 microservices. If 2 microservices do not trust each other, the "OAuth2 client credentials grant" provides a good way to handle the authentication between these 2 microservices [84]. 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 security points taught us that developers should be worried about the security of communication between a microservice and its database (or even other databases) [87]. As an essential part of the data protection strategy, Microsoft strongly advises protecting data in transit [95]. Moreover, because the data is exchanged from many locations, Secure

Sockets Layer (SSL) or TLS protocols are highly recommended. 63 respondents admitted this practice (48.65% *absolutely useful* and 36.49% *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)

👍 “TLS will ensure that **all traffic is encrypted**.” (Technical Lead)

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 2 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 [89].

72.98% of the proponents agreed **PM3** is (*absolutely useful*). Two participants provided conditions for the usefulness of this practice in the following comments.

👍 “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)

👍 “This is useful **when each microservice parses the token, gets the requester information and use that for some operation**. This strongly verifies that the owner of the session is doing that particular operation, rather than reading some parameters to identify who is the owner of the operation.” (Technical Lead)

💡 **PM4.** It is recommended to use the gRPC framework for internal microservice-to-microservice synchronous communication. gRPC is a communication protocol using HTTP2

[96] and Protocol Buffers [97]. The analysis of the security points indicated that gRPC provides an effective solution for direct synchronous communication between microservices [90], [91]. 43 respondents agreed that **PM4** is a (*absolutely useful*) practice (58.11%). More than 25% of the participants were not familiar 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 2 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 [80]. 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. 50 out of 74 survey respondents (67.57%) considered this practice *absolutely useful* or *useful*. We also did not receive *absolutely not useful* for this practice. Two participants shared their 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)

**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

| Security practices for private microservices |   | Sources | AU    | U     | NU    | ANU  | IDK   | MED  | AVG  |
|--|---|---------|-------|-------|-------|------|-------|------|------|
| <b>PP1</b>                                   | Remote nodes (e.g., remote microservices) should not be able to even list/check for the existence of private microservices.                           | [80]    | 40.54 | 27.03 | 13.51 | 0.00 | 18.92 | 3.50 | 3.33 |
| <b>PP2</b>                                   | Use “service grouping” to limit the visibility and callability (of both actions and events) of the private microservices in a group of microservices. | [80]    | 32.43 | 33.78 | 6.76  | 0.00 | 27.03 | 3    | 3.35 |

**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

| Security practices for database and environments |   | Sources          | AU    | U     | NU    | ANU   | IDK   | MED | AVG  |
|--|---|------------------|-------|-------|-------|-------|-------|-----|------|
| <b>PD1</b>                                       | Although security policies should be applied in both development and production environments, production environments need stronger security.   | [82], [98], [94] | 68.92 | 20.27 | 5.41  | 2.70  | 2.70  | 4   | 3.6  |
| <b>PD2</b>                                       | In a microservices architecture, databases should not be exposed to any unauthenticated request.  | [99]             | 71.63 | 18.92 | 4.05  | 2.70  | 2.70  | 4   | 3.64 |
| <b>PD3</b>                                       | 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.  | [100]            | 10.82 | 24.32 | 40.54 | 13.51 | 10.81 | 2   | 2.36 |
| <b>PD4</b>                                       | 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. | [100]            | 18.92 | 35.14 | 24.32 | 4.05  | 17.57 | 3   | 2.84 |

🔗 **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 (49 out of 74) 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 (27.03%). However, **PP2** did not receive *absolutely not useful*, and only a few participants (6.76%) rated it as *not useful*.

An architect commented that service grouping is a beneficial method in microservices systems and is always required. A technical lead also detailed that using this practice can lead to preventing the awareness of systems and improving security.

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

👍 “Ensuring visibility is restricted to the necessary services is **an essential part of zero trust**. It prevents the awareness of systems, **improving security**.” (Technical Lead)

#### 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 (89.19%) 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 [99]. We received a high rate of usefulness (*absolutely useful* or *useful*) for **PD2** (90.55%) which more than 70% 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 65% 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, 35.14% 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 55% of the survey participants agreed with this practice, and less than 30% of our survey participants voted **PD4** as (*absolutely*) *not useful*. Two positive comments on this practice are:

👍 “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)

👍 “One of the problems working with microservices is **data replication**. This should not be a problem if you guarantee that your data will always be updated. You can **use saga pattern or event out box pattern** for this.” (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, 3.5 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 8 most important security practices, including **PA4**, **PI3**, **PI4**, **PI5**, **PI6**, **PM2**, **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 8 highly accepted 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. **PM2** from the “microservices communications” group shows that using security protocols for the communication between microservices and databases is important. **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.

**Several factors still matter while using security practices.** 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. According to the responses of our participants, such factors and trade-offs are enormous, ranging from design context to user experience, from required security skills and expertise to infrastructure resources. For example, it is important to consider if the given system or service is public or private. What are the impacts of the security practices on other quality attributes (e.g., performance)? How sensitive is the data? It is also important to consider to what extent a security practice may impact the user experience. Cost and complexity associated with some practices were mentioned by several respondents as other important factors. A Technical Lead with more than 5 years of experience in microservices systems summarized it as: “*All scenarios [practices] are useful, depending on system require-*

ments. 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**, **PA6**). However, it is out of the scope of our study to explore and discuss all factors and trade-offs. More efforts should be allocated to investigate the short-term and long-term impacts of each of the identified practices, their impacts on a specific type of microservices systems, e.g., IoT microservices systems, and their associated costs and overhead.

***Study why some practices are controversial.*** There are still some practices that were slightly controversial (e.g., **PA1**, **PA2**, **PA3**, **PT4**, **PI1**, **PI2**, **PI3** and **PD1**). For example, **PT4** suggests decoding JWT at the microservices level instead of the API Gateway level (i.e., more than 60% agreed with this practice). At the same time, 25.68% of the practitioners still thought that decoding JWT at the API Gateway level is better than at the microservices level. Also, the survey respondents tended to disagree with a few security practices (e.g., **PD3** and **PD4**) that we found from GitHub and Stack Overflow. There might be several reasons behind controversial practices. As we discussed before, several factors (e.g., the design context and user experience) may impact the opinion of practitioners on using or not using security practices in a software system. Given that the security practices were detected in open-source projects and the survey respondents came from both open-source projects and industrial projects, they might have had a different experience in implementing these practices. Finally, the security practices were provided to the survey respondents in 1-2 sentences, which might not be the best way to describe all aspects of some practices. This can be another reason to cause some controversial practices. Thereby, we argue that an important research direction in future could be exploring the reasons behind controversial practices in securing microservices systems.

***Pay attention to security practices in other or less explored aspects of microservices systems.*** In this study, we only looked at 2 developer discussion platforms (10 microservices systems from GitHub and 306 posts from Stack Overflow) to identify 28 security practices categorized into 6 groups. We do emphasise that these 28 practices are only a subset of available and required security practices for microservices systems. We believe that many security practices were not discussed or could not be found in our data sources. For example, we were not able to find any concrete practices and guidelines regarding security audits, how to safely recover

from security failures and how to secure data in microservices systems. Potential security risks associated with tools and technologies used in microservices system development and deployment also play an important role in achieving secure microservices systems. For example, container images generated by third parties (e.g., Docker) may be associated with several security risks. However, our analysis of open-source projects hosted on GitHub and Stack Overflow posts did not find any concrete practices on how to address the security risks. While some other works (e.g., [7, 10]) 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<sup>6</sup>) 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 [101].

### 6.1. External Validity

Three threats might limit the generalizability of our findings. First, we identified the 28 security practices from only 2 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 software 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. Our validation study did not receive many responses (i.e., 63), and not all the respondents answered the open-ended questions in the survey. 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

<sup>6</sup><https://www.reddit.com>

suggested some feedback. Finally, we deployed a pilot survey to seek practitioners' feedback on the identified security practices. This helped us remove 4 practices and improved the wording of some of the security practices.

The validation survey tried to recruit practitioners with experience in securing microservices systems. As discussed in Section 3.2.2, we used different recruitment approaches for this purpose. One of the approaches was to carefully check the profiles of many practitioners on their websites, Slack, LinkedIn, etc. This approach might have led to two threats. We might have mainly recruited practitioners who successfully applied security in microservices system development and ignored the entire microservices practitioners (e.g., unsuccessful practitioners in applying security in microservices systems). This bias is referred to as survivorship bias [102]. On the other hand, still practitioners with poor knowledge of the MSA style and security may have participated in the survey, which can be a concern for the validity of the survey.

We employed some solutions to (partially) mitigate these threats. Our survey was not filled out by only a few specific roles and did not recruit the respondents using only one recruitment approach. Practitioners with different roles who have worked on various open-source and industrial projects, such as developers, software engineers, DevOps engineers, requirements engineers, and architects, completed the validation survey. We also added the "I Don't Know" option in the survey questions to minimize the concern of lack of knowledge on some identified practices. 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. We invited the contributors of the 10 open-source projects from which the some of best practices were extracted to complete the survey. We acknowledge that the survey responses coming from the contributors of the 10 open-source projects might have provided a biased assessment of the security practices. However, given that we used different recruitment approaches, the percentage of survey responses coming from the contributors of the 10 open-source projects should not be high.

### 6.3. Construct Validity

Our decision to use DeepM1 introduced in [6] 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 GitHub issues and Stack Overflow posts concerning security 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 5 security paragraphs. We admit that some issues or posts with less than 5 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, 9 official documents, and 3 wiki pages from 10 open-source microservices systems, and 306 Stack Overflow posts concerning security in microservices systems. These 28 security practices are categorized into 6 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 74 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.

## Acknowledgements

This work is funded by the National Natural Science Foundation of China (NSFC) with Grant No. 62172311 and the Special Fund of Hubei LuoJia Laboratory.

## References

- [1] N. Dragoni, S. Giallorenzo, A. L. Lafuente, M. Mazzara, F. Montesi, R. Mustafin, L. Safina, Microservices: yesterday, today, and tomorrow, in: *Present and Ulterior Software Engineering*, Springer, 2017, pp. 195–216.
- [2] M. Fowler, J. Lewis, Microservices a definition of this new architectural term (2014).  
URL <https://bit.ly/3zk5xXr>
- [3] P. Jamshidi, C. Pahl, N. C. Mendonça, J. Lewis, S. Tilkov, Microservices: The journey so far and challenges ahead, *IEEE Software* 35 (3) (2018) 24–35.
- [4] P. Di Francesco, P. Lago, I. Malavolta, Architecting with microservices: A systematic mapping study, *Journal of Systems and Software* 150 (2019) 77–97.



- [5] M. Waseem, P. Liang, M. Shahin, A. Ahmad, 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.
- [6] A. Rezaei Nasab, M. Shahin, P. Liang, M. E. Basiri, S. A. H. Raviz, H. Khalajzadeh, M. Waseem, A. Naseri, Automated identification of security discussions in microservices systems: Industrial surveys and experiments, *Journal of Systems and Software* 181 (2021) 111046.
- [7] A. Pereira-Vale, E. B. Fernandez, R. Monge, H. Astudillo, G. Márquez, Security in microservice-based systems: A multivocal literature review, *Computers & Security* 103 (2021) 102200.
- [8] J. Soldani, D. A. Tamburri, W.-J. V. D. Heuvel, The pains and gains of microservices: A systematic grey literature review, *Journal of Systems and Software* 146 (2018) 215–232.
- [9] T. Yarygina, 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.
- [10] A. Hannousse, S. Yahiouche, Securing microservices and microservice architectures: A systematic mapping study, *Computer Science Review* 41 (2021) 100415.
- [11] N. C. Mendonça, C. Box, C. Manolache, L. Ryan, The monolith strikes back: Why istio migrated from microservices to a monolithic architecture, *IEEE Software* 38 (5) (2021) 17–22.
- [12] V. Lenarduzzi, F. Lomio, N. Saarimäki, D. Taibi, Does migrating a monolithic system to microservices decrease the technical debt?, *Journal of Systems and Software* 169 (2020) 110710.
- [13] J. Ghofrani, D. Lübke, Challenges of microservices architecture: A survey on the state of the practice, in: *Proceedings of the 10th Central European Workshop on Services and their Composition (ZEUS)*, CEUR-WS.org, 2018, pp. 1–8.
- [14] O. Zimmermann, Microservices tenets, *Computer Science-Research and Development* 32 (3) (2017) 301–310.
- [15] A. Pereira-Vale, G. Márquez, H. Astudillo, E. B. Fernandez, Security mechanisms used in microservices-based systems: a systematic mapping, in: *Proceedings of the 45th Latin American Computing Conference (CLEI)*, IEEE, 2019, pp. 1–10.
- [16] I. Nadareishvili, R. Mitra, M. McLarty, M. Amundsen, Microservice Architecture: Aligning Principles, Practices, and Culture, O'Reilly Media, Inc., 2016.
- [17] M.-O. Pahl, F.-X. Aubet, S. Liebold, Graph-based IoT microservice security, in: *Proceedings of the IEEE/IFIP Network Operations and Management Symposium (NOMS)*, IEEE, 2018, pp. 1–3.
- [18] M. Waseem, P. Liang, M. Shahin, A systematic mapping study on microservices architecture in devops, *Journal of Systems and Software* 170 (2020) 110798.
- [19] G. Moore, *Crossing the Chasm: Marketing and Selling Technology Project*, HarperCollins Publishers, 2009.
- [20] R. Mahdavi-Hezaveh, J. Dremann, L. Williams, Software development with feature toggles: practices used by practitioners, *Empirical Software Engineering* 26 (1) (2021) 1–33.
- [21] A. Rezaei Nasab, M. Shahin, S. A. Hoseyni Raviz, P. Liang, A. Mashmool, V. Lenarduzzi, *dataset* (2022).  
URL <https://doi.org/10.5281/zenodo.5791337>
- [22] C. Pahl, P. Jamshidi, O. Zimmermann, Architectural principles for cloud software, *ACM Transactions on Internet Technology* 18 (2) (2018) Article No.: 17.
- [23] A. Balalaie, A. Heydarnoori, P. Jamshidi, Microservices architecture enables devops: Migration to a cloud-native architecture, *IEEE Software* 33 (3) (2016) 42–52.
- [24] F. Auer, V. Lenarduzzi, M. Felderer, D. Taibi, From monolithic systems to microservices: an assessment framework, *Information and Software Technology* 137 (2021) 106600.
- [25] D. Taibi, V. Lenarduzzi, C. Pahl, Processes, motivations, and issues for migrating to microservices architectures: An empirical investigation, *IEEE Cloud Computing* 4 (5) (2017) 22–32.
- [26] M. Cinque, R. Della Corte, A. Pecchia, Microservices monitoring with event logs and black box execution tracing, *IEEE Transactions on Services Computing* 15 (1) (2022) 294–307.
- [27] V. Heorhiadi, S. Rajagopalan, H. Jamjoom, M. K. Reiter, 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.
- [28] M. Waseem, P. Liang, M. Shahin, A. Di Salle, G. Márquez, Design, monitoring, and testing of microservices systems: The practitioners' perspective, *Journal of Systems and Software* 182 (2021) 111061.
- [29] J. Bogner, J. Fritzsche, S. Wagner, 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.
- [30] R. Matulevičius, *Fundamentals of Secure System Modelling*, Springer, 2017.
- [31] H. Washizaki, T. Xia, N. Kamata, Y. Fukazawa, H. Kanuka, T. Kato, M. Yoshino, T. Okubo, S. Ogata, H. Kaiya, et al., Systematic literature review of security pattern research, *Information* 12 (1) (2021) 36.
- [32] A. V. Uzunov, E. B. Fernandez, K. Falkner, Assessing and improving the quality of security methodologies for distributed systems, *Journal of Software: Evolution and Process* 30 (11) (2018) e1980.
- [33] J. A. Scott, *A Practical Guide to Microservices and Containers*, MapR Data Technologies, 2018.
- [34] K. Torkura, M. I. H. Sukmana, 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.
- [35] F. Ponce, J. Soldani, H. Astudillo, A. Brogi, Smells and refactorings for microservices security: A multivocal literature review, *arXiv preprint arXiv:2104.13303* (2021).
- [36] P. Billawa, A. B. Tukaram, N. E. D. Ferreyra, J.-P. Steghöfer, R. Scandariato, G. Simhandl, Security of microservice applications: A practitioners' perspective on challenges and best practices, *arXiv preprint arXiv:2202.01612* (2022).
- [37] C. Richardson, *Microservices Patterns: With Examples in Java*, Simon and Schuster, 2018.
- [38] M.-O. Pahl, L. Donini, Securing IoT microservices with certificates, in: *Proceedings of the IEEE/IFIP Network Operations and Management Symposium (NOMS)*, IEEE, 2018, pp. 1–5.
- [39] D. Yu, Y. Jin, Y. Zhang, X. Zheng, A survey on security issues in services communication of microservices-enabled fog applications, *Concurrency and Computation: Practice and Experience* 31 (22) (2019) e4436.
- [40] N. Chondamrongkul, J. Sun, 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.
- [41] Y. Sun, S. Nanda, 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.
- [42] T. Bi, P. Liang, A. Tang, X. Xia, Mining architecture tactics and quality attributes knowledge in stack overflow, *Journal of Systems and Software* 180 (2021) 111005.
- [43] I. Malavolta, G. A. Lewis, B. Schmerl, P. Lago, D. Garlan, Mining guidelines for architecting robotics software, *Journal of Systems and Software* 178 (2021) 110969.
- [44] N. Meng, S. Nagy, D. Yao, W. Zhuang, 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.
- [45] B. G. Glaser, A. L. Strauss, E. Strutzel, The discovery of grounded theory: strategies for qualitative research, *Nursing Research* 17 (4) (1968) 364.
- [46] B. A. Kitchenham, S. L. Pfleeger, *Personal opinion surveys*, in: *Guide to Advanced Empirical Software Engineering*, Springer, 2008, pp. 63–92.



- [47] S. O. member, Single sign-on in microservice architecture (July 2021).  
URL <https://stackoverflow.com/questions/25595492>
- [48] G. member, Authorization between services (July 2021).  
URL <https://github.com/moleculerjs/moleculer/issues/304>
- [49] G. member, Identity/customer service as a microservice (July 2021).  
URL <https://github.com/dotnet-architecture/eShopOnContainers/issues/785>
- [50] G. member, Single sign on: Azure ad b2c vs identityserver4, and others (July 2021).  
URL <https://github.com/dotnet-architecture/eShopOnContainers/issues/925>
- [51] S. O. member, How to refresh request token with microservice multiple instances? (July 2021).  
URL <https://stackoverflow.com/questions/48760583>
- [52] G. member, Discussion on security (July 2021).  
URL <https://github.com/spinnaker/spinnaker/issues/148>
- [53] S. O. member, Should api gateway be responsible for authorisation? (June 2018).  
URL <https://stackoverflow.com/questions/50700178>
- [54] G. member, Securing ui of jaeger (June 2017).  
URL <https://github.com/jaegertracing/jaeger/issues/218>
- [55] S. O. member, Micro-service architecture, should the spring cloud config server, zuul gateway server and eureka server be protected as resources? (July 2021).  
URL <https://stackoverflow.com/questions/61307668>
- [56] G. member, Proposal: Create the template function for authentication in the file for each service (July 2021).  
URL <https://github.com/goadesign/goa/issues/2361>
- [57] S. O. member, How to authenticate json web tokens (jwt) across different apis? (July 2021).  
URL <https://stackoverflow.com/questions/55394912>
- [58] G. member, gateway api (July 2021).  
URL <https://github.com/dotnet-architecture/eShopOnContainers/issues/239>
- [59] S. O. member, Laravel passport, oauth and microservices (July 2021).  
URL <https://stackoverflow.com/questions/39126827>
- [60] G. member, Rfc: Allow spring property placeholders in pipeline expressions (August 2019).  
URL <https://github.com/spinnaker/spinnaker/issues/4725>
- [61] S. O. member, Should jwt be a separate auth micro-service and not sit with the backend business logic? (July 2021).  
URL <https://stackoverflow.com/questions/58948497>
- [62] G. member, Startup.cs - add authorization with ocelot (July 2021).  
URL <https://github.com/dotnet-architecture/eShopOnContainers/issues/647>
- [63] S. O. member, In a microservice environment, should any producer be able to verify jwt tokens? (July 2021).  
URL <https://stackoverflow.com/questions/44471051>
- [64] S. O. member, Shared signature key for jwt in various microservices (July 2021).  
URL <https://stackoverflow.com/questions/43559197>
- [65] G. member, Rfc: Halyard secret management (November 2018).  
URL <https://github.com/spinnaker/spinnaker/issues/3649>
- [66] G. member, Hide passwords in urls on the /config endpoint (February 2020).  
URL <https://github.com/cortexproject/cortex/pull/2176>
- [67] G. member, Vulnerable data exposed with metrics endpoint (March 2019).  
URL <https://github.com/jaegertracing/jaeger/issues/1428>
- [68] G. member, grpc plugin framework does not respect -query.bearer-token-propagation flag (September 2019).  
URL <https://github.com/jaegertracing/jaeger/issues/1821>
- [69] G. member, Cortex feature request/improvement - refresh aws object store credentials for expired tokens (January 2021).  
URL <https://github.com/cortexproject/cortex/issues/3731>
- [70] S. O. member, Microservices - how to solve security and user authentication? (July 2021).  
URL <https://stackoverflow.com/questions/32574103>
- [71] S. O. member, Decoding oauth2 jwt at api gateway level vs at individual microservice level (July 2021).  
URL <https://stackoverflow.com/questions/51524648>
- [72] G. member, Find the best location to inject server information to the routing handler (July 2021).  
URL <https://github.com/networknt/light-4j/issues/11>
- [73] G. member, Add logging module for light 4j rfc#29 (July 2021).  
URL <https://github.com/networknt/light-4j/issues/453>
- [74] JWT, Introduction to json web tokens.  
URL <https://jwt.io/introduction>
- [75] G. member, How is https/ssl termination handled? (July 2021).  
URL <https://github.com/dotnet-architecture/eShopOnContainers/issues/752>
- [76] S. O. member, 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? (July 2021).  
URL <https://stackoverflow.com/questions/54631411>
- [77] G. member, Iam authentication support in ruler and alertmanager s3 client (August 2020).  
URL <https://github.com/cortexproject/cortex/issues/3034>
- [78] G. member, Authenticating to gcp when using chunks storage (bigtable and gcs) (October 2020).  
URL <https://github.com/cortexproject/cortex/issues/3306>
- [79] S. O. member, Oauth2 grant for server-to-server communication (July 2021).  
URL <https://stackoverflow.com/questions/27838280>
- [80] G. member, Private services (July 2021).  
URL <https://github.com/moleculerjs/moleculer/issues/124>
- [81] S. O. member, How to authenticate and authorize in a microservice architecture? (July 2021).  
URL <https://stackoverflow.com/questions/52349986>
- [82] G. member, Addtocart method relies on the posted productdetails (July 2021).  
URL <https://github.com/dotnet-architecture/eShopOnContainers/issues/250>
- [83] S. O. member, How to add an api with oauth2 on the top of kong? (July 2021).  
URL <https://stackoverflow.com/questions/47324184>
- [84] S. O. member, Quick solution to handle service to service authentication in a microservices architecture (July 2021).  
URL <https://stackoverflow.com/questions/61433192>
- [85] G. member, Span authentication support in jaeger collector (September 2017).  
URL <https://github.com/jaegertracing/jaeger/issues/427>
- [86] G. member, Flaky test: Testreload (November 2020).  
URL <https://github.com/jaegertracing/jaeger/issues/2622>
- [87] S. O. member, Microservices and database security (July 2021).  
URL <https://stackoverflow.com/questions/53621693>
- [88] G. member, Add tls client reload (August 2020).  
URL <https://github.com/cortexproject/cortex/issues/3012>
- [89] S. O. member, Oauth 2.0 in microservices: When a resource server communicates with another resource server (July 2021).  
URL <https://stackoverflow.com/questions/52290697>
- [90] G. member, eshoponcontainers (July 2021).  
URL <https://github.com/dotnet-architecture/eShopOnContainers/wiki/gRPC>
- [91] J. member, Jaeger (July 2021).  
URL <https://www.jaegertracing.io/docs/1.27/architecture/>
- [92] G. member, Jaeger trace sampling should not be decided by every service (by default) (July 2018).  
URL <https://github.com/cortexproject/cortex/issues/885>
- [93] G. member, Build a secure channel for security reports (October 2017).  
URL <https://github.com/jaegertracing/jaeger/issues/457>
- [94] G. member, Allow secure communication between components (October 2017).  
URL <https://github.com/jaegertracing/jaeger/issues/458>

- [95] M. member, Azure data security and encryption best practices (2021).  
URL <https://docs.microsoft.com/en-us/azure/security/fundamentals/data-encryption-best-practices#protect-data-in-transit>
- [96] I. Grigorik, Surma, Http/2.  
URL <https://developers.google.com/web/fundamentals/performance/http2>
- [97] Google, Protocol buffers.  
URL <https://developers.google.com/protocol-buffers>
- [98] G. member, Deploying spinnaker with halyard to k8s with kube v2 provider and ssl enabled for gate fails because k8s readinessprobe fails (May 2018).  
URL <https://github.com/spinnaker/spinnaker/issues/2765>
- [99] S. O. member, Login authentication flow for microservices (July 2021).  
URL <https://stackoverflow.com/questions/59058573>
- [100] G. member, After customerbasket has been posted to basketcontroller where is the unitprice validated with the catalog in the workflow? (July 2021).  
URL <https://github.com/dotnet-architecture/eShopOnContainers/issues/945>
- [101] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, A. Wesslén, Experimentation in Software Engineering, Springer, 2012.
- [102] S. J. Brown, W. Goetzmann, R. G. Ibbotson, S. A. Ross, Survivorship bias in performance studies, The Review of Financial Studies 5 (4) (1992) 553–580.