A Survey on Microservices Security: Preliminary Findings

Davide Berardi[1], Saverio Giallorenzo[2], Jacopo Mauro[2],
Andrea Melis[1], and Fabrizio Montesi[2]

[1] Università di Bologna
[2] University of Southern Denmark

1 Introduction and Methodology

In recent years Microservices have become the state-of-the-art architectural style for distributed systems [7]. Despite its widespread adoption—and possibly due to its recent introduction—we notice the lack of comprehensive guidelines on Microservices Security. Motivated by this observation, we started an ongoing Systematic Literature Review process to categorise the literature on Microservices Security, with the intent to overview the current status of the field, to provide an initial guideline to researchers and practitioners, and to possibly identify uncovered areas and orient future research.

Our presentation In our presentation at Microservices 2020, we propose to present the state of the art and discuss some of the most interesting insights we collected so far, such as: *a*) main security attacks occurred, *b*) the relation between software development and security, *c*) the main technologies involved, and *d*) the main application domains. Among the findings we will present, there are a set of categorisations of our dataset that help in describing the current shape of the research field (e.g., aggregation points and "uncovered" areas). As an example, one categorisation we performed is the creation of a word-net of semantic relations (reported in section 2) to cluster the dataset around the concepts that characterise each publication. We will use that and similar categorisations to structure the presentation of our findings. An example of insight emerged from our study is that more than a third of the papers in our dataset pair microservices security with the usage of the DevOps development methodology [3]. This reflects the recent attention gathered by a variant of DevOps, centred around site reliability engineering [4] and security, called DevSecOps [5, 13]. In section 2 we briefly report the main take aways from the above findings. We conclude this section with a brief account of the methodology we followed to collect our dataset.

Methodology To collect our dataset, we searched with the—purposefully generic—query "Microservice AND Security" over the 5 main libraries of peer-review literature in the field of Computer Science: ACM Digital Library [1], IEEE Explore [10], SpringerLink [14], Scopus [9], Wiley [15] and ScienceDirect [8]. This first search returned a total of 539 papers. On these papers, we applied a rejection process based on the disjunction of the following criteria: *R0*) the paper is published before 2015; *R1*) security is not a topic of the paper; *R2*) microservices are not a topic of the paper; *R3*) the work does not include a detailed contribution (e.g., it is a position paper, research proposal, or poster); *R4*) the paper added contributions on microservice security but was out of our scope: microservices topic was ortogonal or the security topic was incidental on the main contibutions of the paper. After the filtering phase, which totalled 168 papers, we followed a fixed-point backward snowballing process [11] with the rejection criteria above—each new entry in the dataset triggering a new snowballing phase—totalling 184 papers

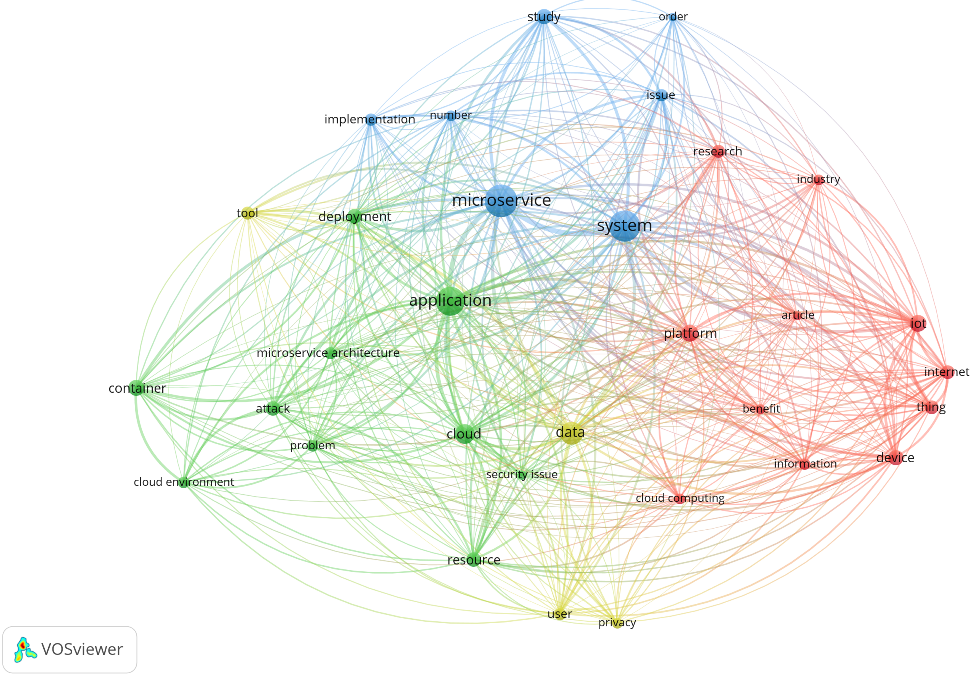

2 Preliminary Findings

Wordnet We generated a word-net, shown in the figure above, based on the text of each paper, identifying several clusters to analyse the main context area of our dataset.

In the figure, the different colours indicate four clusters. The green cluster corresponds to *technical terms* where the main focus goes on the technology, such as microservice architecture, container, deployment, and resource, related to security (security issue, attack, problem). The red cluster identifies the *settings* (research, industry) and the *application areas* of the paper (internet-of-things, iot, cloud computing). The blue cluster corresponds to *concepts*, like microservice, system, issue, and study. The yellow cluster identifies the *subjects of research*, like user, privacy and data. Notably, tool, although belonging to this cluster, is a bridge-word between the green and blue clusters, i.e., a link between the technical and the conceptual realms.

DevSecOps Around 38% of papers from our dataset link DevOps and Microservices in the context of security. DevOps is a recent software development methodology that combines software development (Dev) and information-technology operations (Ops) into "a set of practices intended to reduce the time between committing a change to a system and the change being placed into normal production, while ensuring high quality" [3]. Since 2018 we notice the growth of the concept of DevSecOps [12, 2], an extension of DevOps that advocates the integration of security practices into the DevOps approach. Some relevant examples are [2, 6], where the authors have studied and implemented methods to integrate the certification of software components within continuous integration.

References

[1] ACM. ACM Digital Library. https://dl.acm.org, Acc. May 2020.

[2] M. Anisetti, C. A. Ardagna, F. Gaudenzi, and E. Damiani. A continuous certification methodology for devops. In *Proceedings of the 11th International Conference on Management of Digital EcoSystems*, pages 205–212, 2019.

[3] L. Bass, I. Weber, and L. Zhu. *DevOps: A software architect's perspective*. Addison-Wesley Professional, 2015.

[4] B. Beyer, C. Jones, J. Petoff, and N. R. Murphy. *Site Reliability Engineering: How Google Runs Production Systems*. O'Reilly Media, Inc., 2016.

[5] J. Bird. Devopssec: Securing software through continuous delivery. 2016.

[6] G. Casale, M. Artač, W.-J. van den Heuvel, A. Van Hoorn, P. Jakovits, F. Leymann, M. Long, V. Papanikolaou, D. Presenza, A. Russo, et al. Radon: rational decomposition and orchestration for serverless computing. *SICS Software-Intensive Cyber-Physical Systems*, pages 1–11, 2019.

[7] N. Dragoni, S. Giallorenzo, A. Lluch-Lafuente, M. Mazzara, F. Montesi, R. Mustafin, and L. Safina. Microservices: Yesterday, today, and tomorrow. In M. Mazzara and B. Meyer, editors, *Present and Ulterior Software Engineering*, pages 195–216. Springer, 2017.

[8] Elsevier. ScienceDirect. https://www.sciencedirect.com/, Acc. May 2020.

[9] Elsevier. Scopus. https://www.scopus.com, Acc. May 2020.

[10] IEEE. IEEE explore. https://ieeexplore.ieee.org, Acc. May 2020.

[11] S. Jalali and C. Wohlin. Systematic literature studies: database searches vs. backward snowballing. In *Proceedings of the 2012 ACM-IEEE international symposium on empirical software engineering and measurement*, pages 29–38. IEEE, 2012.

[12] S. Mansfield-Devine. Devops: finding room for security. *Network Security*, 2018(7):15–20, 2018.

[13] H. Myrbakken and R. Colomo-Palacios. Devsecops: a multivocal literature review. In *International Conference on Software Process Improvement and Capability Determination*, pages 17–29. Springer, 2017.

[14] Springer. SpringerLink. https://link.springer.com, Acc. May 2020.

[15] Wiley. Wiley Editor. https://onlinelibrary.wiley.com, Acc. May 2020.