Towards Securing the 6G Transition: A Comprehensive Empirical Method to Analyze Threats in O-RAN Environments

Felix Klement, *Graduate Student Member, IEEE*, Wuhao Liu and Stefan Katzenbeisser, *Senior Member, IEEE*

***Abstract*—In this paper, we present a new methodology that enables the MITRE ATT&CK framework to objectively assess specific threats in 6G Radio Access Networks (RANs). This helps address new security challenges that arise in the transition to open RANs. We analyze the O-Cloud component within the O-RAN ecosystem as a representative example, wherein no individual threat class demonstrates complete security. The inherent modularity of our approach ensures great adaptability and allows it to be applied to various other components within this system. This allows us to effectively detect and combat threats, thereby ensuring the resilience and security of future communication networks.**

***Index Terms*—Security, Open RAN, Telecommunication, MITRE ATT&CK, CVE, CWE.**

I. Introduction

THE ongoing progression from 5G to 6G networks represents a significant paradigm shift in communications technology, leading to groundbreaking advancements. As this transition continues alongside the simultaneous evolution of Open RAN methodologies, it becomes crucial to confront the emergence of various new threats that present substantial obstacles to network security. Working Group (WG) 11 of the O-RAN Alliance has conducted research to identify several potential threats, which demand thorough analysis and effective mitigation strategies to guarantee the security and robustness of forthcoming networks. Nevertheless, the present analysis exhibits numerous deficiencies, relying on subjective and conventional methodologies for risk assessment, thereby underscoring the imperative for enhancements. Although initial investigations have been undertaken in the realm of risk and threat analysis pertaining to emerging telecommunication systems ([1], [2], [3]), no singular method has been established that possesses both reusability and universal applicability.

This paper introduces a novel methodology that combines the utilization of the MITRE ATT&CK framework with empirical data to evaluate specific threats during the transition towards open 6G networks. Our study concentrates on evaluating threats within the O-RAN ecosystem, with a particular emphasis on the O-Cloud component as a representative example. This is a cloud-native computing platform comprising a cluster of physical infrastructure nodes that encompasses all pertinent O-RAN elements. Consequently, it offers enhanced adaptability and scalability for RAN provisioning, alongside various other benefits. However, the modular nature of our approach allows for its application to all other components as well. The primary objective is to detect and analyze potential attack surfaces or vulnerable targets that malicious actors can exploit.

Our approach encompasses a rating system that provides a comprehensive assessment of the vulnerability levels exhibited by diverse technologies and platforms in a given scenario. This scoring mechanism facilitates a rapid evaluation of the most severe and perilous threats by aggregating baseline scores. Moreover, we introduce more granular scoring metrics by employing average scores from the Common Vulnerability Scoring System (CVSS), thereby enhancing the scoring capacity beyond the conventional three score classes employed by the O-RAN Alliance (*High, Medium and Low*). To identify potential countermeasures against the identified threats, we leverage established resources such as the MITRE techniques, Common Weakness Enumeration (CWE), and Common Vulnerabilities and Exposures (CVE). By extracting countermeasures from these resources, we strengthen the ability to effectively address and mitigate the identified threats. Our investigation reveals notable exploitability score associated with the O-Cloud platform, in contrast to previous traditional RAN deployment methods. This fact can be attributed to its extensive attack surface resulting from its inherently adaptable characteristics. Additionally, we identify that four specific threats within the realm of credential access and authentication exhibit a substantial range of potential tactics. Moreover, we observe that the graphical depiction that we create for the cumulative base scores offers a straightforward and efficient approach to effectively administer and structure prioritization endeavors concerning vulnerability remediation.

In summary, the objective of this work is to make a contribution to the ongoing research endeavors focused on enhancing the security of the 6G transition. This is achieved through a comprehensive analysis of the distinct threats present in O-RAN environments.

Our approach integrates empirical data, the MITRE ATT&CK framework, and robust assessment methodologies to gain a profound understanding of vulnerabilities and potential countermeasures specific to the individual RAN com-

Manuscript received 28 December, 2022; revised 2 June, 2023; accepted 2 August, 2023. The authors acknowledge the financial support by the German Federal Ministry of Education and Research – Bundesministerium für Bildung und Forschung (BMBF), as part of the Project "6G-RIC: The 6G Research and Innovation Cluster" (project number 825026). *(Corresponding author: Felix Klement.)*

Felix Klement, Wuhao Liu and Stefan Katzenbeisser are with the University of Passau, Faculty of Computer Science and Mathematics, Innstraße 43, 94032 Passau, Germany (e-mail: felix.klement@uni-passau.de; wuhao.liu@uni-passau.de; stefan.katzenbeisser@uni-passau.de)

ponents. By effectively identifying and addressing the associated threats, we can strategically deviate from potential risks. Through an enhanced comprehension of the individual vulnerabilities and the implementation of appropriate defensive measures, we can ensure the resilience and security of forthcoming communication networks.

The structure of this paper is as follows: Section II describes relevant work that is within the scope of our research. In Section III, we provide an overview of the O-Cloud. Then, in Section IV, we outline the methodology of our approach. The results of our analysis are presented and evaluated in Section V. We then briefly describe opportunities for future improvements to our approach in Section VI. Finally, in Section VII, we summarize our main findings and conclusions.

II. Related Work

The field of risk and threat analysis is extensively documented in previously published works and studies ([4], [5], [6], [7], [8] and [9]), frequently leveraging the MITRE ATT&CK framework [10]. However, these analyses tend to focus on classical computing systems or software development, with limited attention paid to mobile communication systems. Initial efforts have been made to adapt the framework for use in mobile communication systems, with the development of MITRE FiGHT [11] serving as a platform that empowers cybersecurity professionals to leverage up-to-date threat data, tools, and techniques for integration into their security operations. Nonetheless, given the increasing frequency of security breaches in mobile networks, the need for effective threat modeling and risk management strategies in this area is pressing. Although there exist publications in this domain as we describe in Section II-B, they frequently suffer from excessive specialization or limited universal applicability, rendering them unsuitable for automated replication. Thus, we propose extending existing mechanisms to make them applicable and more usable in the context of O-RAN.

A. Threat modeling in mobile communication systems

In [2], Chen et al. construct a theoretical framework for threat modeling grounded in the MITRE ATT&CK framework, which provides a structured classification of malicious behaviors in end-to-end mobile communications. They also executed a user study involving industry experts to evaluate and investigate its potential applications. The findings demonstrate the potential advantages of such a framework and associated tools for the mobile communications industry. Regrettably, the framework is not publicly accessible, rendering external scientific evaluation and potential expansion challenging. This underscores a deficiency of openly accessible tools in the domain of telecom security research.

The study conducted by Pell et al. [12] aims to enhance protection against Advanced Persistent Threats (APTs) in 5G networks by addressing the gaps in the current 5G threat assessments and the MITRE ATT&CK threat modeling framework. The authors identified certain areas of knowledge deficiency in the present framework, specifically regarding crucial 5G technology enablers like Software Defined Network (SDN), Network Function Virtualization (NFV), and 5G-specific core network signaling protocols. By analyzing previous attacks on telecommunication networks and the intentions of APT groups, the authors demonstrated how APTs could exploit domain-specific techniques in multi-stage attack scenarios. The study recommends an approach that can support a comprehensive cyber risk assessment, intrusion detection, and the development of protective measures for 5G core networks. The findings of this research have been incorporated into MITRE FiGHT, with close collaboration between MITRE and Pell, including the implementation of some of Pell's suggestions.

In the paper by Wang et al. [13], a method for building a knowledge graph of cyber attack behavior based on Common Attack Pattern Enumerations (CAPECs) and CWEs to support network intelligence in 6G networks is presented. The knowledge graph implemented in the Neo4j graph database can be used to comprehensively capture the strategies and behaviors of specific attacks and provide clues for attack prediction and network situational awareness in 6G networks. Overall, the paper argues that a knowledge graph based on the CAPECs and CWEs databases can be a useful tool for improving security in 6G networks.

B. Open RAN specific investigations

Moreover, scant literature exists regarding the evaluation of threats and risks associated with the O-RAN methodology. Nevertheless, a handful of recent publications have emerged elucidating the realm of security in Open RAN. It is crucial to differentiate between Open RAN, a comprehensive designation for decentralized systems featuring open and compatible interfaces, and O-RAN, the specific progression towards such an architectural framework. Researchers in [14] have devised a taxonomy and an extensive survey that encompasses various categories of risks prevalent in the domain of Open RAN. The authors discern a range of attack vectors, encompassing but not limited to fronthaul assaults, breaches compromising data integrity and confidentiality, compromises in monitoring mechanisms, component integrity infringements, and unauthorized physical access to components. While the article examines security best practices for Open RAN, it falls short of subjectively assessing individual risks. In this context, our proposed methodology can be employed as a complementary approach to address this gap effectively.

Polese et al. [15] conduct a thorough examination of the O-RAN specification, its architectural components, and their operational aspects. They explore security-critical elements and propose potential strategies for enhancing the security of RAN deployments. While the study briefly touches upon relevant threats such as those targeting the O-Cloud, there is an absence of direct quantification regarding the extent of individual security issues within the system's context, along with a corresponding deficiency in their accurate evaluation.

In [16] they present a comprehensive examination of the evolutionary trajectory of the Open RAN. The paper offers a thorough analysis of the constituent technologies utilized within the Open RAN framework, accompanied by a discussion of the corresponding projects, activities, and standardization efforts. Furthermore, the paper addresses the prevalent

challenges encountered in the development and implementation of Open RAN, while identifying future research directions. The authors emphasize the potential advantages of Open RAN, underscore its significance, and elucidate the obstacles associated with its realization. The scholarly insights provided in this paper serve as a foundational reference and launchpad for further investigations in this field.

The authors in [3] discuss the Open RAN architecture, which aims to promote innovation and competition in the RAN market. They analyzes the components of Open RAN deployments, evaluate their security state, and propose measures for secure operation. The paper highlights the advantages of Open RAN, clarifies its relationship with O-RAN and OpenRAN, and explains its architectural breakdown. It emphasizes the need for early security assessment and discusses the incorporation of security concepts in the design phase. The paper concludes that Open RAN does not introduce major security issues and suggests defining security methodologies for critical points in such deployments.

The presented literature provides evidence of the efficacy of employing the MITRE ATT&CK framework for threat modeling within mobile communication systems. Additionally, recent publications on security in Open RAN highlight a deficiency in empirical evaluation and assessment of key areas in telecommunication networks. This knowledge gap serves as the starting point for our proposed approach, aiming to address the significant gaps and shortcomings in evaluating specific RAN components. To bridge this gap, we propose a novel modular approach that integrates MITRE ATT&CK and its associated tools (CAPEC, CWE, CVE).

III. The O-Cloud Platform

In the ensuing discourse, we provide a concise overview of the O-Cloud, a cloud computing platform delineated by the O-RAN Alliance. It encompasses a constellation of physical infrastructure nodes that house diverse prerequisites and pertinent functionalities, alongside accompanying software components and supporting management and orchestration mechanisms. We also present the diversity of O-Cloud deployments and the challenges associated with them. The majority of our data is derived from the specifications put forth by the Alliance ([17], [18], and [19]). Our objective is to furnish the reader with a comprehensive understanding of the component and its inherent susceptibilities within the specified framework.

A. Introduction to the O-Cloud

The O-Cloud is a cloud-based computing platform that includes a collection of physical infrastructure nodes and hosts the components of relevant O-RAN functions (e.g., Near-Real-Time RAN Intelligent Controller (Near-RT RIC), O-RAN Central Unit (O-CU-CP), O-RAN Central Unit - Control Plane (O-CU-CP), O-RAN Central Unit - User Plane (O-CU-UP), O-RAN Distrubuted Unit (O-DU), O-RAN Radio Unit (O-RU) logical functions). Beyond that, it also provides support for software components and the corresponding management and orchestration functions. Generally, it can be seen as the

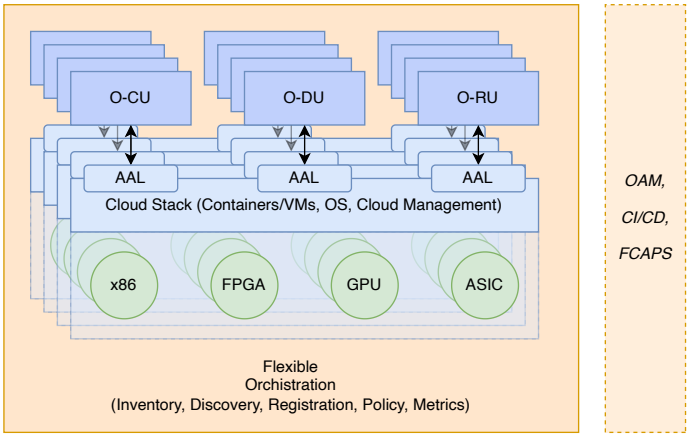

Fig. 1. Essential components for the orchestration and cloudification

central execution environment for virtualized O-RAN components [19]. This enables greater flexibility and scalability in the deployment of RANs, making it easier to integrate new technologies and innovations into the network, to better meet their customers' changing requirements.

A simplified schematic representation of the individual components within a potential O-Cloud deployment is shown in Figure 1. A so-called node within the deployment consists of a variety of different hardware components (Central Processing Unit (CPU), memory, disk space, etc.). Furthermore, O-Cloud components can encapsulate additional technologies for acceleration. An example of such a technology would be Field-Programmable Gate Arrays (FPGAs) to perform certain computations and thus gain a significant performance advantage over running on a traditional CPU. Other hardware accelerators are Graphics Processing Units (GPUs), Digital Signal Processings (DSPs), Application-Specific Integrated Circuits (ASICs) but also acceleration functions like Forward Error Correction (FEC), Low-Density Parity-Check (LDPC), Artificial Intelligence (AI) or specific security algorithms. The respective setup is operated by the individually defined cloud stack. This can be realized in many different ways: One example would be a deployment using OpenStack and Kubernetes on Commercial Off-the-Shelf (COTS) hardware, interconnected by a spine/leaf networking fabric. Moreover, the inclusion of Operations, Administration and Maintenance (OAM) tools, alongside the implementation of Continuous Integration / Continuous Deployment (CI/CD) practices, is essential for facilitating seamless operations and streamlined administration. Additionally, it is highly advisable to align all processes with the Fault, Configuration, Accounting, Performance and Security (FCAPS) model to the greatest extent feasible. FCAPS, following the ISO model for telecommunication network management, encompasses fault, configuration, accounting, performance, and security management tasks, thereby guaranteeing the integrity and security of system components. The orchestration of all these instances and components is defined using the Accelerator Deployment Model (ADM), or more specifically via the Acceleration Abstraction Layer (AAL) interface. The respective manufacturer of an acceleration component must provide an open interface to

ensure seamless integration into the system. The O-RAN Alliance cites open Application Programming Interfaces (APIs) such as Data Plane Development Kits (DPDKs) CryptoDev, EventDev and Base Band Device (BBDEV) [17] as already known examples.

B. Deployment Variations

With regard to the O-Cloud, there are a large number of possible deployment variations. In its current specification, based on the location of deployment, three types of cloud infrastructures are defined: Regional Cloud (RC-D), Edge Cloud (EC-D) and Cell Site (CS-D). Among them, CS-D refers to the location where Physical Network Functions (PNFs) such as O-RU is, and the other two are where Cloudified Network Functions (Cloudified NFs) are located. These can be combined as desired out of different use cases and requirements. A primary important influencing factor is the latency requirement between O-Cloud functions. For example, O-CU-CP can be deployed on RC-D with Near-RT RIC or on EC-D with O-DU because of different latency limitations.

In [17] the various possible combinations of the different variants are defined. Figure 2 shows how the individual network functions (in the upper part of the figure) can be used and combined either as a PNF or as a Cloudified NF. This results in a total of six different scenarios, each of which differs according to the use of the respective O-RAN Key Technology (ORKT) and the deployment location used (RC-D, EC-D, CS-D). The term *O-Cloud* within the colored boxes in Figure 2 refers to the fact that the RAN functions used are supported by an O-RAN Cloud platform. O-RAN PNF means that it is a full-fledged physical O-RAN network function. There may also be hybrids of the two where appropriate. We will not go into more detail about the respective scenarios in this paper, these are described in depth by [17].

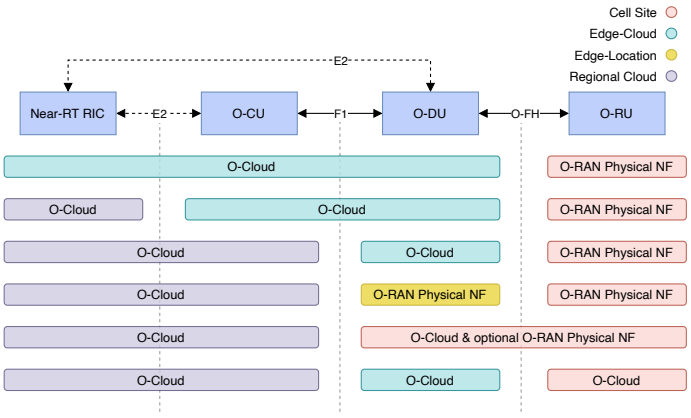

Fig. 2. High-level comparison of potential deployment scenarios

Besides the specific O-Cloud deployment scenarios, general cloud deployment modes have to be considered. According to National Institute of Standards and Technology (NIST) SP 800-145[1], they are: private cloud, community cloud, public cloud and hybrid cloud. The difference between them is the degree of control over the underlying resources. The private cloud allows operators access exclusively to all the physical resources, while the public cloud makes the resources available to the public through services controlled by the cloud providers. It is conceivable that the level of control over resources is directly proportional to the investment in security maintenance.

[1]https://csrc.nist.gov/publications/detail/sp/800-145/final

C. New Challenges through Cloud Environments

The shift of RAN components into a cloud environment automatically results in a completely new landscape. The O-Cloud interface incorporate a range of innovative technologies that have the potential to significantly expand the attack surface of a deployment. The variety and novelty of many of the technologies used, especially the combination of virtual technologies and the telecommunication domain, have not yet been sufficiently investigated from a security perspective, creating additional vulnerabilities that can be exploited by malicious actors. One of the fundamental distinctions that must be differentiated is the following: the cloud infrastructure and computing resources are managed exclusively by a single operator or managed using commercially operated cloud solutions offered by service provider (AWS, Google Cloud etc.), namely, private cloud or public cloud. This makes the final threat surface to be considered enormously different. In most recently published risk analyses ([1], [3]) it was stated that cloud providers hosting O-Cloud components theoretically have the same capabilities as traditional RAN operators. In other words, in the case of public cloud, a malicious cloud provider could compromise the security of a RAN. This could be mitigated to some extent by deploying private cloud. In this case, the O-Cloud can be accorded the same level of confidentiality as the operator of the RAN itself. However, the private cloud increases other cost aspects, such as maintenance and upgrades.

To sum up this section, the combination of RAN and cloud, along with the diversity of the deployment increases the risk of possible security vulnerabilities, which can, however, be reduced to a minimum by means of well-structured countermeasures. It is therefore all the more important that there is a uniform and publicly accessible concept for managing and analyzing threats within these O-Cloud deployments. A major problem we have encountered with regard to all-around security within the O-Cloud is the following: The objective of the O-RAN Alliance is solely focused on ensuring that its specifications serve as guiding principles for implementation purposes. Therefore, there are only a few binding security measures at the current time. In order to ensure the protection of a deployment against a potentially malicious cloud provider, we recommend the rigorous implementation of security measures. These measures should include secure access protocols and binding security requirements, which should be applied in addition to the existing O-Cloud specifications. By following these recommendations and implementing a on-premise deployment, it is anticipated that the risk of security breaches in both virtualized and cloud-based RANs will be minimized. A first approach towards such a solution is explained in Section IV.

IV. Methodology

To perform a empirical in-depth analysis of each O-Cloud threat, we employ the MITRE ATT&CK framework. The following section is divided into four parts. First, we provide a brief overview of the framework's structure and the tools we use to create our initial dataset, including the meaning of the individual sources. After that, we briefly go over how exactly the mapping of a threat to a technique works. We then discuss the assumptions and relationships we established to create the dataset as well as gain valuable insights from it. Finally, we describe the used libraries for the implementation of the empirical evaluations.

A. MITRE ATT&CK Framework

The MITRE ATT&CK framework is a comprehensive knowledge base of adversary Tactics, Techniques and Procedures (TTPs) based on real-world observations [20]. It includes detailed information about the tactics, techniques, and sub-techniques used by attackers, as well as examples, mitigations, and detection methods. The framework is organized in a matrix that maps the tactics and techniques used by attackers to the phases in their respective attack campaigns. Thus, it is possible to identify which tactics and techniques are used together to potentially successfully achieve the attack objective. As of today, a total of 14 tactics with over 400 techniques have been captured. These are regularly updated based on observations from the field. The framework has become a widely adopted standard for both understanding and defending against cyberattacks.

In addition to the points just listed, there are other reasons why O-Cloud security analysis should be combined with MITRE ATT&CK. One reason is that the framework is a useful instrument for holistic system security assessment through the use of CAPECs, CWEs as well as CVEs. Those tools are maintained by the non-profit organization MITRE. Together, these three resources are used by security professionals to help identify and prevent potential attacks on computer systems. The framework also provides a more detailed and comprehensive examination of individual identified problem areas using information from the previously mentioned tools. Furthermore, it is open-source and community-based, which means that continuous feedback can be obtained from the cybersecurity community.

B. Threat to Technique Mapping

In order to effectively utilize the tactics and techniques outlined in the MITRE framework, we align the O-Cloud threats identified by WG 11 of the O-RAN Alliance [21]. To achieve this objective, we employ suitable methodologies within the matrix specifically designed for cloud technologies. This approach is the sole manual allocation factor in our methodology. However, this is shown to be feasible due to the thorough description of the individual techniques and their areas of application. In future scenarios, envisioning the simplification of this process is plausible through the utilization of advanced computational models, such as large-scale language models. By means of semantic comparison, we assess the extent of similarity between O-RAN threats and attack techniques. While certain cases exhibit conspicuous keywords that indicate a high degree of relevance between a technique and a threat (e.g., the threat "Build image on VL" and the technique "Build image on Host"), most instances are more intricate. The semantic similarities lie beneath the surface of the text and require meticulous examination. For instance, the threat "Abuse a O-Cloud administration service" and the technique "Container Orchestration Job" may not exhibit significant similarity in their titles alone. However, upon closer inspection, phrases such as "Adversaries may abuse a container administration service to execute commands within a container" and "adversaries may abuse task scheduling functionality provided by container orchestration tools such as Kubernetes to schedule deployment of containers configured to execute malicious code" reveal noteworthy relevance. This discerning semantic analysis facilitates the establishment of a comprehensive mapping between techniques and threats, which is documented in Table I. Our analysis reveals that, while most of the O-Cloud threats could be matched with techniques from the MITRE matrix, there are some threats that are specific to the architecture of the O-RAN system and therefore can not be fully aligned with common cloud applications. In the future, is plausible that a dedicated matrix will be developed, similar to the work conducted by Pell et al. in creating an adapted matrix for 5G networks [12], or the ongoing development of MITRE FiGHT, to enhance the precision of security evaluations for O-RAN systems. Nevertheless, for the purpose of our initial proof-of-concept, we consider this limitation to be acceptable.

The findings of our attribution analysis are presented in Table I. This tabular representation enumerates the methodologies employed for each threat, as indicated in the MITRE column. These assigned methodologies, derived from the ATT&CK framework, serve as the foundation for computing comprehensive scores for each unique threat identifier.

C. Metrics Score Calculation

The Base Score Metrics (BSMs) are defined as a vector, as shown in Equation 1. Individual metrics can be derived from this vector to calculate different scores for a CVE. In this investigation, we solely focus on the BSM of a CVE, as the availability of Temporal Score Metrics and Environmental Score Metrics is often limited for numerous CVEs.

$$\begin{aligned} AV:\{L;A;N\} = AC:\{H;M;L\} = Au:\{M;S;N\} \\ C:\{N;P;C\} = I:\{N;P;C\} = A:\{N;P;C\} \end{aligned} \tag{1}$$

The vector is segmented into six metrics. The Access Vector (AV) depicts the attacker's potential entry points into the system. It is further classified into Local (L), Adjacent Network (A), and Network (N). Access Complexity (AC) evaluates the level of difficulty for an attacker to exploit a vulnerability. AC is categorized as High (H), Medium (M), or Low (L). Authentication (AU) determines whether the attacker requires authentication to exploit a vulnerability. AU is segmented into Multiple (M), Single (S), or None (N). The degree of impact on Confidentiality (C), Integrity (I), and Availability (A) of

TABLE I
LISTING OF THREATS DEFINED BY [21] WITH THE ADDED AVERAGED CVSS INFORMATION

Defined by the O-RAN Alliance [21]							Avg. CVSS			MITRE
Cat.	Thread-ID	Description	CIA	Severity	Likelihood	Risk Score	Impact	Exploitability	Base Score	Assigned ATT&CK Technique
Generic	T-GEN-01	Software flaw attack	C,I	●	●	●	2.9	10.0	5.0	T1068
	T-GEN-02	Malicious access to exposed services using valid accounts	C,I	◉	●	●	6.1	9.7	7.2	T1078
	T-GEN-03	Untrust binding between the different O-Cloud layers	C,I	●	●	●	-	-	-	/
	T-GEN-04	Lack of Authentication & Authorization in interfaces between O-Cloud components	C,A	●	●	●	-	-	-	/
	T-ADMIN- 01	Denial of service against NFO/FOCOM	A	◉	◉	●	4.3	9.2	5.6	T1498
	T-ADMIN- 02	Abuse a O-Cloud administration service	C,I,A	◉	●	●	5.4	8.5	6.1	T1552, T1609, T1204
Virtual-Machines/Containers	T-VM-C-01	Abuse of a privileged VM/Container	C,I,A	●	●	●	3.4	9.5	5.1	T1016
	T-VM-C-02	VM/Container escape attack	C,I,A	●	◉	●	3.4	9.5	5.1	T1611, T1538
	T-VM-C-03	VM/Container data theft	C,I	●	●	●	3.6	8.2	4.8	T1530, T1552, T1609, T1538
	T-VM-C-04	VM/Container migration attacks	C,I,A	●	●	●	4.2	9.1	5.4	T1040, T1609, T1499, T1496
	T-VM-C-05	Changing virtualization resource without authorization	A	●	◉	●	4.3	9.3	5.6	T1578, T1499
	T-VM-C-06	Failed or incomplete VNF/CNF termination or releasing of resources	C	◉	◉	◉	-	-	-	
Images	T-IMG-01	VM/Container images tampering	C,I	◉	●	◉	6.4	7.7	6.4	T1195, T1525, T1610, T1612, T1204
	T-IMG-02	Insecure channels with images repository	C,I	◉	●	●	4.1	8.2	5.0	T1040
	T-IMG-03	Secrets disclosure in VM/Container images	C,I	◉	◉	◉	5.8	6.9	5.6	T1600, T1195, T1552,
	T-IMG-04	Build image on VL	C,I,A	◉	◉	◉	6.7	6.3	5.9	T1612
	T-VL-01	VM/Container hyperjacking attack	C,I,A	●	●	●	4.0	9.5	5.5	T1036, T1068, T1496
	T-VL-02	Boot tampering	I	●	◉	◉	6.4	9.2	7.1	T1542, T1495
Interfaces	T-O2-01	MitM attacks on O2 interface between O- Cloud and Service Management and Orchestration (SMO)	C,I,A	○	●	●	3.8	7.9	4.6	T1613, T1040, T1082, T1580, T1070, T1609, T1049, T1619, T1046
	T-OCAPI- 01	MitM attacks on O-Cloud interface between VNFs/CNFs and the virtualization layer	C,I,A	○	●	●	4.1	8.2	5.0	T1040
Resources	T-HW-01	Cross VM/Container side channel attacks	C,I,A	●	●	◉	6.2	9.5	7.1	T1003, T1204, T1614
	T-HW-02	MitM attacks on the interface between virtualization layer and hardware	C,I,A	●	●	●	-	-	-	/

O-RAN Scores: ● ≡ High, ◉ ≡ Medium, ○ ≡ Low — Avgerage CVSS: $10 \cdots 6.67 \equiv$ High, $6.\overline{6} \cdots 3.34 \equiv$ Medium, $3.\overline{3} \cdots 0 \equiv$ Low

information stored in the system is described by the respective metrics. The degree of each impact is classified as None (N), Partial (P), or Complete (C). The evaluation assigns a score of 0.0 to 10.0 to each metric. The individual values are calculated using the standardized CVSS [22]. The Equations 2 to 5 are used to calculate the scores used in our evaluations.

$$\begin{aligned} ConfImp &= 1 \quad ConfImpact \\ IntImp &= 1 \quad IntegImpact \\ AvailImp &= 1 \quad AvailImpact \\ &= ConfImp \quad IntImp \quad AvailImp \\ ) \; Impact &= 10{:}41 \quad (1 \quad) \end{aligned} \tag{2}$$

$$\begin{aligned} AccComp &= AccessComplexity \\ Auth &= Authentication \\ AccVec &= AccessVector \\ &= AccComp \quad Auth \quad AccVec \\ ) \; Exploitability &= 20 \end{aligned} \tag{3}$$

$$f(Impact) = \begin{cases} 0 & \text{if } Impact = 0 \\ 1{:}176 & \text{otherwise} \end{cases} \tag{4}$$

$$\begin{aligned} Imp &= 0{:}6 \quad Impact \\ Exp &= 0{:}4 \quad Exploitability \\ &= Imp + Exp \quad 1{:}5 \\ ) \; BaseScore &= f(Impact) \end{aligned} \tag{5}$$

As an illustrative case, let us consider the threat T-GEN-01, which is an identifier designated by the O-RAN Alliance and pertains to software vulnerabilities resulting in attacks. When a successful attack occurs, both confidentiality and integrity are compromised. The severity, likelihood, and risk score values provided by the Alliance are all categorized as *High*, as indicated in the corresponding column. To establish a correlation to the MITRE ATT&CK framework, we have utilized our mapping process (refer to Section IV-B) and assigned the technique T1068. Subsequently, we employ a python script to extract the associated CWEs. Within this context, we have identified three specific CWEs (CWE-552, CWE-706, CWE-46), leading to the programmatic assignment of in total ten CVEs. We then retrieve the BSM vectors for each of these CVEs. Then, for each CVE, the impact, exploitability, and base score value is determined. As an illustrative example, we calculate these scores for CVE-2001-0693 which has the following vector:

$$AV:L = AC:L = Au:NC:P = I:P = A:P \tag{6}$$

from where the corresponding numerical values for each component of the vector are derived, as observed in Equation 7.

$$\begin{aligned} AV:0{:}395 = AC:0{:}71 = Au:0{:}704 \\ C:0{:}275 = I:0{:}275 = A:0{:}275 \end{aligned} \tag{7}$$

Subsequently, the values obtained can be substituted into the formulas (Equations 2 to 5), thereby yielding distinct outcomes for each respective CVE. Equation 8 provides the numerical values associated with CVE-2001-0693. This procedure is repeated for all ten CVEs, followed by the computation of their average values. Consequently, these averaged results are obtained and correspond to the values presented in the respective column of Table I.

$$\begin{aligned} Impact &= 6{:}44297677 \\ Exploitability &= 3{:}948736 \\ BaseScore &= 4{;}63964982 \end{aligned} \tag{8}$$

D. Dataset Composition and Structure

The first step in creating our dataset, is to find CAPECs over all the techniques defined for each O-Cloud threat class. Subsequently, all of them are identified. And in the last procurement step, associated CVEs are searched for the CWEs found. Generally speaking, the relationship between the individual tools could be represented as in Equation 9. This shows the interrelationships between the different tools.

$$CVE \quad CWE \quad CAPEC \tag{9}$$

The use of CVEs is due to the robust metrics implemented by the CVE standard. These metrics, such as the *Access Vector*, *Access Complexity*, and *Authentication* capture the invariant characteristics of a vulnerability in terms of time and user environment, allowing for accurate representation of the associated hazard. The inclusion of impact metrics, which evaluate the direct impact on the confidentiality, integrity, and availability of an asset, further enhances the utility of CVEs. Compared to CWE and CAPEC, which only offer likelihood and severity (only for CAPECs) ratings, CVEs provide a more comprehensive assessment of vulnerabilities.

In Equation 5, the resulting final score represents the intrinsic characteristics of a vulnerability that remain constant over time and across different user environments. However, using only the base score to assess risk can sometimes lead to incorrect assessments. Therefore, it is often more effective to conduct a more comprehensive risk assessment by considering contextual factors and additional attributes. In our dataset, we also store values such as the exploitability score, impact score, and access vector to account for these factors. Furthermore, our use of the JavaScript Object Notation (JSON) schema to represent individual data allows for easy extensibility of the dataset to include additional considerations outside of CVSS.

The respective CVSS values obtained by applying Equation 2, 3 and 5 for a CVE together with further information like the *v2_vector* (version 2 of the cvss vector as described in Equation 1), the *access_vector* (access strategy derived form the cvss vector) are stored. In addition to these values, the complete metrics provided by the API, as well as other relevant metadata, are also included. A sample of the a potential resulting dataset is shown Listing 1.

In total we identify 46 CAPEC categories for 47 techniques. To determine the CVEs for each CAPEC category, we isolate the associated CWE for each category. In order to eliminate duplicates, we first check the list of CAPECs for duplicates and remove any that are found. By storing the direct assignment of each CAPEC to its corresponding technique in

advance, we are able to avoid unnecessary duplicate queries. At the end of this process, we have a total of 109 CWEs. Finally, we can query the list of CVEs for each individual CWE and thereby obtain a total of 781 CVEs, which are then stored in a list corresponding to the CWE.

At the moment we only use the entities defined in the framework. In the future, it is conceivable that extensions based on MITRE ATT&CK (such as [12]) will also be added to the resource pool of possible techniques to include an even larger set of possible vulnerabilities.

```
"scan_date": "2022-12-09",
"scan_runtime": "00h 46m and 27.11s",
"data": [
  {
    "technique_id": "T1498",
    "t_findings": [
      {
        "capec_id": "CAPEC-125",
        "c_findings": [
          {
            "cwe": "CWE-404",
            "cves": [
              {
                "id": "CVE-1999-1127",
                "score": [...],
                "v2_score": 5,
                "v2_exploitability_score": 10,
                "v2_impact_score": 2.9,
                "v2_vector": "",
                "access_vector": "",
                "full_metrics": [...],
                "description": "",
                "cpe_vulnerable": true/false,
                "cpe_criteria": "",
                "published": "1999-12-31T05:00:00.000",
                "last_modified": "2018-10-12T21
                    :29:22.827"
              },
              ...
            ],
            "cwe_info": {...}
          },
          ...
      },
    ...
  ]
```

Listing 1. Example of a simplified possible JSON dataset

E. Implementation

Our codebase consists of two jupyter notebooks: one dedicated to data gathering, which generates a dataset in the JSON format, and another for analysis, wherein evaluations are conducted. Our approach heavily relies on five open-source libraries. For handling ATT&CK content, we utilize the MITRE-provided python implementation[2]. To establish links between CAPEC entries and their corresponding techniques, we require a means of searching for techniques stored in Structured Threat Information Expression (STIX) 2 format. Thus, we employ the python API[3] provided by the OASIS Technical Committee for serializing and deserializing STIX2 JSON content. For querying CWEs in Python, the cwe2[4] library is employed. To obtain information on CVEs, we utilize both the cve_lookup[5] library and NVDLib[6], which serves as a convenient wrapper around the. All of the code and datasets generated throughout our research are publicly available on GitHub[7].

V. Results

To enhance the analysis of the gathered data, we generate multiple visual representations. These visualizations serve as illustrative instances for users employing our open-source methodology, granting them the freedom to assess the datasets formatted with JSON using their preferred techniques. Moreover, in conjunction with our published source code, we supply a collection of python helper functions that facilitate graphical assessments. This empowers users to create diverse visualizations of significance to them in the future and subsequently assess them. We provide a succinct summary of the key insights derived from the graphs we generated to demonstrate the applicability of our empirical data collection from Sections V-B to V-D.

A. O-Cloud CVSS Threat Scores

The risk assessment procedure outlined in [21] by the Security WG 11 of the O-RAN Alliance encompasses several key steps. These include the identification of assets, threats, and vulnerabilities, followed by an evaluation of the associated risk. The assessment of threat criticality is based on the potential severity of its consequences and the likelihood of its occurrence. Severity is gauged by considering factors such as the impact on data protection, confidentiality, integrity, and availability, as well as the extent of resource affected and the efficacy of existing controls. Severity levels are classified into three categories: low, medium, and high, with low denoting the least severe and high representing the most severe. The risk value is subsequently determined using the RISK formula ($Severity \quad Likelihood$). Nonetheless, it is important to note that this risk determination approach is not without its limitations, which we shall elucidate upon briefly. Consequently, we will provide a comprehensive explanation of the calculated scores and expound upon their respective significance.

Traditional risk assessment methods such as those used in [21] have a notable drawback in their dependence on subjective severity ratings. These ratings introduce subjectivity and allow for varying interpretations, thereby resulting in inconsistencies within the risk assessment process. To mitigate this issue, we have devised an alternative approach that leverages the MITRE ATT&CK framework. This framework offers a systematic ranking of emerging vulnerabilities grounded in empirical observations of adversary behavior, thereby circumventing subjective assessments. Consequently, our approach provides a consistent and objective foundation for comprehending and analyzing cyber threats. By adopting this methodology, our evaluation offers a comprehensive and objective overview of the likelihood and impact of specific threats, thereby eschewing the reliance on subjective assessments.

Moreover, the risk assessment model employed by the Alliance exhibits a deficiency in terms of adaptability, potentially rendering it less effective in dynamically evolving

[2]https://github.com/mitre-attack/mitreattack-python
[3]https://github.com/oasis-open/cti-python-stix2
[4]https://github.com/nexB/cwe2
[5]https://github.com/MachineThing/cve_lookup
[6]https://github.com/Vehemont/nvdlib/
[7]https://github.com/fklement/acema_oran

circumstances or when confronted with new information. This rigidity restricts its ability to remain relevant over time. In contrast, our evaluation methodology offers the advantage of iterative execution, enabling swift generation of results that can promptly adapt to changing circumstances and incorporate newly acquired information.

Lastly, the severity and likelihood model employed in conventional risk assessment methods may exhibit excessive complexity or necessitate specialized expertise, rendering it challenging for certain individuals to comprehend and implement. In contrast, our evaluation methodology offers a streamlined, objective approach to assess and prioritize vulnerabilities consistently. This approach serves as a valuable tool for enhancing the effective management of cybersecurity risks.

Having elucidated the issues and limitations encountered in the assessment process when employing the conventional methodology, we now proceed to present the outcomes attained through the utilization of our novel approach for the computation of CVSS scores. The application of the methodologies outlined in our approach enables the derivation of the overall average Base, Impact, and Exploitability scores for the O-Cloud component, resulting in the following scores:

O-Cloud Base Score: 5.8
O-Cloud Impact Score: 4.7
O-Cloud Exploitability Score: 9.0

The utilization of averaged scores proves instrumental in the identification, monitoring, and mitigation of risk within a given system or a set of interconnected systems. By consistently computing and juxtaposing averaged risk scores, it becomes feasible to discern variations in risk and make well-informed decisions regarding risk management strategies moving forward. As a general observation, it can presently be affirmed that the O-Cloud component stands at a heightened probability of being subjected to a successful attack, attributable to its exceptionally elevated exploitability score.

For the threats where a mapping is currently possible, we calculate the respective average CVSS scores and assign them in Table I. Using the CVSS, we obtain the impact, exploitability, and associated base score value. However, these metrics cannot be compared one-to-one with the severity, probability, and risk values defined by the O-RAN Alliance. Impact refers to the extent to which something affects or changes something else, while severity refers to the intensity or severity of the problem. For example, in the context of a risk assessment, impact may refer to the potential consequences of an event or situation, such as financial loss, injury, or property damage. Severity, on the other hand, may refer to the likelihood that an event or situation will occur and the extent of the impact it could have. In the best case, a risk analysis has both values consulted for evaluation. Impact, however, is generally shown to be more useful because it is more closely related to the consequences or outcomes of an event or situation. It refers to the extent to which something affects or changes something else and is generally considered a key factor in determining the overall effects and potential consequences of an event or situation. This information can be useful in identifying the risks with the greatest potential impact and prioritizing the allocation of resources to address them. Another reason impact is considered more important or useful than severity is that it is more objective and easier to quantify. While severity is often subjective and depends on the perspective of the person assessing the risk, impact can often be measured more objectively. This makes it easier to compare the relative impact of different risks and make informed decisions about how to manage them. If we now consider the impact scores in Table I, we notice that T-IMG-04 has a very high impact value. However, this is not reflected in the severity rating of the O-RAN Alliance and the threat would generally not be considered as significant. The other threats can also be compared much better in terms of impact and thus evaluated by the score.

In the context of a risk assessment, likelihood may refer to the probability that a vulnerability will be exploited or that an undesirable event will occur. Exploitability, on the other hand, refers to the ease with which a vulnerability can be exploited or taken advantage of to cause harm. In other words, exploitability refers to the extent to which an attacker can exploit a vulnerability to gain unauthorized access to a system or cause damage. One reason why we consider this more important or useful than the likelihood is that it helps to identify vulnerabilities that pose the greatest risk to a system. Vulnerabilities with high exploitability are generally considered more risky because they are easier for attackers to exploit, which means that they are more likely to be successfully exploited and to cause harm. By identifying threats with high exploitability, it's possible to prioritize efforts to address and to minimize the risk of an attack. Now, if we examine the exploitability scores in Table I, we notice that all threats except one are classified as high. Again, a more specific distinction is now possible. Thus one would first look at the threats which are in the upper value range of the High CVSS score class (such as T-VM-C-01, T-VL-01, T-GEN-01 etc.).

With the factors we introduced in the O-Cloud risk analysis we provide a more complete picture of the potential risk of a vulnerability or weakness. Taken together, these factors can help identify vulnerabilities that are more likely to be exploited and can have significant impact if exploited. The factors to consider when assessing the risk of a system ultimately depend on the context and information available. It is important to use an approach to risk assessment that is appropriate for the situation and takes into account all relevant factors. However, if we now directly compare the scores in Table I defined by the O-RAN Alliance with those calculated by us, we can see our empirical approach provides more information about the individual threats.

B. Comparison of Threat Base Scores

Figure 3 presents a comparison of individual O-RAN threats and their corresponding techniques based on their accumulated base scores, which have been divided by severity level (*Low, Medium, High*). A bar chart with accumulated risk scores divided by severity can be a valuable tool in security risk analysis as it offers a clear visual representation of the relative risk levels of various security issues. By accumulating the

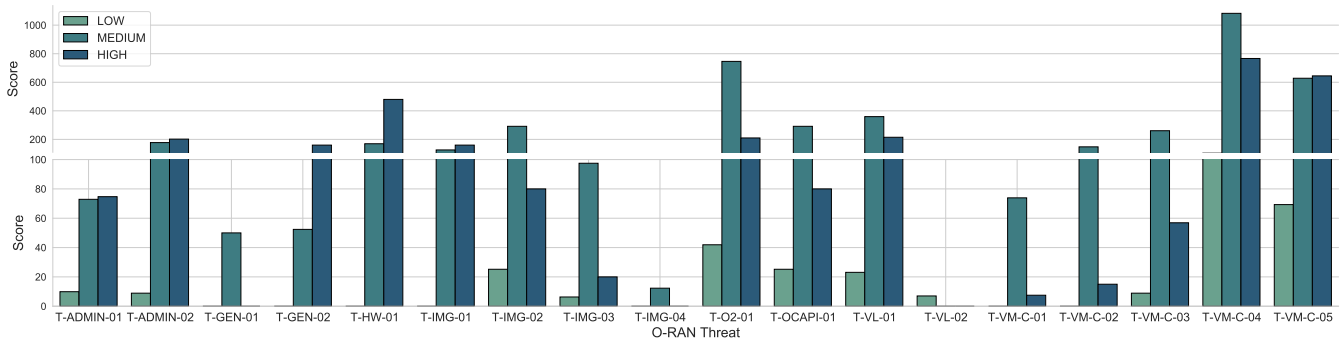

Fig. 3. Accumulated base score ratings divided by severity level per threat

risk scores, the bar chart illustrates the overall risk level for each issue, enabling the operator of an O-Cloud component to quickly identify the most pressing issues that require immediate attention and prioritize threats that should be secured or monitored more closely. In addition, the bar chart allows for convenient comparison of the risk levels of different issues, helping the operator prioritize their efforts and allocate resources efficiently. Overall, the bar chart with accumulated risk scores can aid in the understanding and management of risks in security risk analysis.

The results show that threats related to VM/container migration attacks (T-VM-C-04) and Man-in-the-Middle (MiM) attacks on the O2 interface between the O-Cloud and SMO (T-O2-01) have an average high score of vulnerabilities with medium severity. Notably, only three mapped threats do not have vulnerabilities with high severity. The T-HW-01 threat, which pertains to cross VM/container side channel attacks, presents a significant number of vulnerabilities with high severity, as do T-VM-C-04 and -05. One potential strategy for mitigating this threat would be to implement an on-premise hosting solution for the O-Cloud, as this would eliminate the possibility of collocated VM/containers being targeted by attackers exploiting these vulnerabilities within the O-Cloud.

C. Evaluation of Tactics Coverage

Tactics in the MITRE ATT&CK framework refer to the specific goals or objectives that an adversary is attempting to achieve through their TTPs. Therefore, it is of significant importance to analyze and assess the various O-Cloud threats in terms of the tactics they employ. The tactics in the framework are Initial Access, Execution, Persistence, Privilege Escalation, Defense Evasion, Credential Access, Discovery, Lateral Movement, Collection, Exfiltration, Command and Control, and Impact. Understanding the tactics that an adversary is using can help security professionals to identify and respond to potential threats and to develop strategies for defending against future attacks.

In order to ascertain the tactics that an attacker may use in relation to O-Cloud threats, we analyze the number of identical tactics accumulated for each individual threat using a heat map (see Figure 4). Four threats, T-ADMIN-02, T-HW-01, T-IMG-03, and T-VM-C-03, have a high number of potential tactics in the area of credential access, which can be used to gain access to credentials or other authentication information. The threat T-02-01 shows a presence in the areas of defense evasion and discovery, with defense evasion tactics being employed to evade detection and prevent discovery by security measures and discovery tactics being used to gather information about a target system or network. These threats should be prioritized when investigating a running O-Cloud component and determining its vulnerability to related tactics. To reduce the likelihood of these tactics being successful, our python scripts can be used to access individual techniques and retrieve the associated mitigations and detection mechanisms. This information can be utilized to identify potential threats and formulate strategies to counter future attacks.

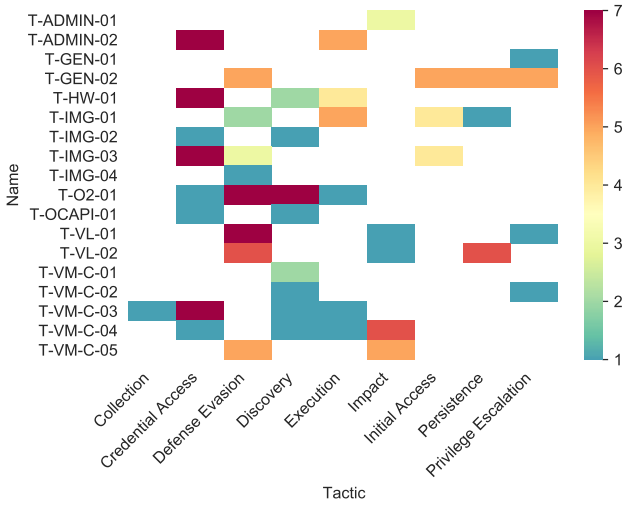

Fig. 4. Heat map of accumulated identical tactics per threat

D. Estimation of Platform Occurencies

The term ”platforms” within the MITRE framework refers to specific operating systems, software, and hardware that are the targets of analysis. The framework currently includes eight main platforms: Windows, Linux, MacOS, Android, iOS, AWS, GCP, and Azure. Additionally, the cloud matrix expands

to include platforms such as Office 365, Google Workspace, containers, Software as a Service (SaaS), and Infrastructure as a Service (IaaS). These platforms encompass a range of systems and devices, and the framework offers a comprehensive overview of the tactics and techniques that adversaries may employ to attack them. It is worth noting that the Cloud matrix is focused specifically on tactics and techniques that are specific to cloud environments. Other matrices in the MITRE ATT&CK framework, such as the Windows and Linux matrices, cover tactics and techniques that may be used to target those specific platforms, regardless of whether they are deployed in a cloud or on-premises environment.

We have counted and sorted the number of individual platform occurrences per technique based on their accumulated values. The results show that Windows is the most vulnerable platform, especially when it comes to threats in the O-Cloud component. This is likely due to its wide distribution, versatility, and large user base, which make it a popular target for attackers. However, it is worth noting that Windows also has a well-established ecosystem of security tools and procedures in place to protect against these threats, as well as a robust security infrastructure with features like user account control and antivirus built into the operating system. Further stands out that IaaS is more vulnerable to attacks in the MITRE ATT&CK framework than SaaS because the user has more control and is responsible for securing the infrastructure, which can be customized and integrated with other systems, increasing the attack surface. With IaaS, the user also has more visibility into the infrastructure and its components, which can be both a benefit and a risk. While this allows the user to better understand and secure the infrastructure, it also gives an attacker more information about the infrastructure and potential vulnerabilities to exploit. Overall, IaaS requires more effort and expertise to secure and can have a larger attack surface compared to SaaS, where the service provider is responsible for securing the infrastructure. It is important to consider these points if the O-Cloud or any of its components are deployed in a model of this type.

VI. Future Work

As previously mentioned, our analysis would greatly benefit from the implementation of a customized MITRE matrix designed specifically for the O-RAN approach. One potential approach to create such a matrix involves adapting the matrix developed for 5G networks, as proposed in Pell's work on modeling the 5G core [12]. This adapted matrix can then be expanded to encompass the Open RAN context. Moreover, the integration of the FiGHT framework into our workflow presents an additional potential avenue for incorporating a tailored matrix specifically designed for 5G networks.

Furthermore, apart from the platforms encompassed within the MITRE ATT&CK framework, there exist supplementary frameworks and taxonomies dedicated to specific system types or technologies. Notably, the Industrial Control Systems Cyber Emergency Response Team (ICS-CERT) framework is tailored for Industrial Control Systems (ICS), while the Taxonomy for Supervisory Control and Data Acquisition (SCADA) systems

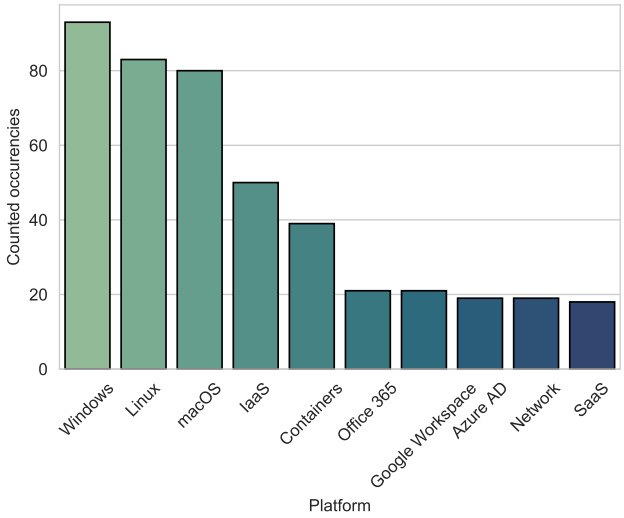

Fig. 5. Comparison of counted platform occurrences in associated MITRE ATT&CK techniques

offers targeted categorization. These frameworks furnish more granular insights into the tactics and techniques employed by malicious actors to exploit these specialized system types. Consequently, an advantageous augmentation to our empirical approach would involve incorporating these frameworks.

In forthcoming research endeavors, we aim to enhance the efficacy of the mapping procedure between O-RAN threats and attack techniques through the utilization of language models. Our intention is to automate and achieve complete quantification of this mapping process.

VII. Conclusion

Our empirical assessment, in conjunction with the MITRE ATT&CK framework, enables a more scientifically grounded and effective analysis of individual threats to the O-Cloud as identified by the O-RAN Alliance. By employing visualization techniques, we facilitate the efficient sorting and management of efforts aimed at ensuring secure operation of the O-Cloud component. This methodology can be consistently applied to diverse datasets, providing a valuable tool for continuously assessing vulnerability levels. We illustrate how this approach can isolate and delve into the most severe potential threats, studying their impact in a comprehensive manner.

Based on our evaluation of the O-Cloud component, within the cumulative base score rating, merely three of the predefined threat classes exhibit an absence of vulnerabilities rated as *High*. Furthermore, there exists no threat class devoid of any vulnerabilities, albeit threat classes T-VL-02 and T-IMG-04 display an exceedingly scant number of cumulative vulnerabilities. Furthermore, the assessment of tactics coverage reveals intriguing findings pertaining to the realm of viable tactics. Specifically, within the domain of credential access, a notable four threats demonstrate a significant proportion of potential exploitable shares.