

Misbehavior Detection and Mitigation on 5G Core Services in Kubernetes

Hristo Koshutanski
*Eviden Strategy & Innovation
R&D Spain
Atos IT Solutions and Services
Iberia
Madrid, Spain
0000-0002-0660-0635*

Sarang Kahvazadeh
*Services as networkS (SaS)
Centre Tecnologic
Telecomunicacions Catalunya
(CTTC)
Barcelona, Spain
0000-0001-5607-8120*

Jesús Villalobos
*Eviden Strategy & Innovation
R&D Spain
Atos IT Solutions and Services
Iberia
Seville, Spain
0000-0002-4703-983X*

Alejandro Garcia Bedoya
*Eviden Strategy & Innovation
R&D Spain
Atos IT Solutions and Services
Iberia
Seville, Spain
0000-0002-7203-5734*

Josep Mangues-Bafalluy
*Services as networkS (SaS)
Centre Tecnologic
Telecomunicacions Catalunya
(CTTC)
Barcelona, Spain
0000-0003-4960-9434*

Abstract— Kubernetes plays an important role nowadays in enabling cloud-native telecom and managing microservices in 5G/6G networks. It provides a platform for deploying, scaling, and managing cloud-native micro-services but also 5G/6G core functionalities. In this work, we demonstrate the importance of misbehavior detection in Kubernetes clusters paired with suitable mitigation measures as a solid baseline to minimize disruption of 5G service delivery. We evidence the impact of two types of DoS attacks, blackholing and network DoS, and results of detection and mitigation in Kubernetes.

Keywords—Decentralized Intrusion Detection, Microservices, 5G/6G, Kubernetes, Security.

I. INTRODUCTION

Container-based microservices management platforms such as Kubernetes offer highly dynamic and automated computing clusters' creation and management, multiple computing clusters co-existence and operation, and modular networking between all containers to enable the final business logic.

Given the increasing dynamicity and complexity of computing clusters nowadays:

- It becomes challenging to identify erroneous behaviors inside clusters,
- It increases vector of attacks and room for malicious activities to take place undetected, and
- It becomes essential to monitor all components to ensure the entire cluster is healthy.

The detection of anomalies in computing clusters has recently gained attention in the literature [1,2,3,4]. The strong point of related work is advancing the knowhow on how AI/ML algorithms facilitate learning and detection of anomalous

behaviors in Kubernetes and Docker environments, and possible deployment settings.

We demonstrate a novel work on network anomaly detection that scales to the dynamic nature of Kubernetes to realize misbehavior detection and ensure the healthy status of clusters with no, or acceptable, deviation from the normal behavior. The work distinguishes from the state-of-the-art on the following:

- Consistent flow telemetry over volatile and dynamic IP addresses.
- Flexible workflow between flow telemetry on each worker node and decentralization of deep learning models on one or more nodes in a cluster.
- Fine-grained association of models to a set of protocols and/or nodes in a cluster for enriched behavior insight and detection,
- Full visibility of network traffic across nodes and virtual NICs in a node,
- Realistic testbed of 5G service provisioning for representative validation.

We use case two types of DoS attacks, blackholing and network DoS, at layers 2 and 3 of the OSI model, and their impact. We demonstrate the importance of anomaly detection and mitigation on 5G core services deployed in Kubernetes.

II. DEMONSTRATION TESTBED

The demonstration is based on a 5G testbed provisioned by CTTC, and the Eviden's technology for anomaly detection [6], named LADS.

A. 5G Provisioning

In CTTC's testbed, a Kubernetes cluster is implemented consisting of one master node and three worker nodes. Figure 1 illustrates the architecture of the testbed. The Kubernetes cluster architecture has been meticulously configured to support both the Open5GS core network and a video-on-demand microservice, utilizing a suite of enabling tools such as MetalLB for load balancing, Prometheus and Grafana for monitoring, and persistent volumes for data storage.

The Open5GS core is deployed within a dedicated namespace (5Gcore-noupf), intentionally excluding the User Plane Function (UPF), which is instantiated separately in its own namespace (UPF). In parallel, the video-on-demand streaming microservice is implemented and deployed within the vstream namespace.

To facilitate end-to-end 5G connectivity, the Open5GS core operating within the Kubernetes environment is routed and integrated with the SrS-RAN platform [5], which is physically deployed at the CTTC premises. This setup enables real-world testing scenarios, leveraging a variety of commercial smartphones.

Furthermore, components for decentralized intrusion detection and mitigation are co-located within the same Kubernetes cluster, providing enhanced security functionalities as part of the overall architecture.

B. Decentralized Intrusion Detection and Mitigation

The notion of intrusion detection is realized by (i) decentralization of a network telemetry sensor on every node in a Kubernetes cluster including the master node, and (ii) decentralization of the machine learning decision making module, called LADS-Brain, on one or more nodes in a cluster. Figure 2 shows the workflow of intrusion detection and mitigation on a set of nodes in a cluster. In this case, one LADS-Brain instance is deployed on one of the nodes and associated to the whole cluster. The sensor offers visibility of network telemetry per node or per segment, while the LADS Brain manages and assigns deep learning models corresponding to a family of protocols in a cluster.

The Mitigator component has two main roles: (i) Monitors the event log of the LADS-Brain and matches a set of rules for each new event added to decide on need of mitigation action. For instance, if a UPF pod or service receives a high-volume incoming traffic, say more than 500% of packet/s (pps) or bytes/s (bps) than those in training; and (ii) Offers a set of named mitigation action and their underlying execution means, such as scripts that interact with the Kubernetes control plane to stop, move, restart pods. In the case of a potential DoS against the UPF service, the Mitigator reallocates the UPF pod to another worker to ensure continuity of the service.

III. ATTACK SCENARIOS

We performed the following network attacks:

- Blackholing through ARP poisoning on the UPF service. Particularly, spoofing the MAC address of the cluster gateway in the UPF pod's ARP table. The aim is to DoS (blackhole) the UPF pod's outgoing external communications, and

consequently the whole 5G service provisioning. As a result, the UPF pod's outgoing communications are directed to the spoofed MAC address of the gateway's IP, which drops all of them. This is an OSI layer 2 cyber-attack. The diagram shows the two main attack instances of blackholing (with just 2 spoofed packets per second attack footprint). Figure 3 shows the traffic footprint of the attack with two visible packet drops of the impact.

- ICMP ping flood on the UPF to show another attack with a similar DoS impact but of a high volumetric nature on the OSI layer 3. We controlled the attack in two main spikes of 20K pps for a duration of 15 seconds each. Figure 4 shows the traffic footprint of the attack performed.

Both types of attacks – Blackholing ARP poisoning and ICMP ping flood, on the UPF service are based on Insecure Workload Configurations allowing unauthorized access to the host of a node, and privilege escalation to execute false packets against the victims' pods on the node.

IV. RESULTS

The results for the different attack instances are promising. In all attack cases, an accuracy over 96% was achieved, with an F1-Score for legitimate cases exceeding 97%, and an F1-Score for anomalous cases on average 90%. In some cases, there were some false positive errors in predicting legitimate traffic as anomalous traffic, but this is an acceptable error. Given that the opposite error, misclassified anomalous traffic as normal traffic, is the most critical for anomaly detection systems. In our experiments, the detection of all attack instances showed no failure (no false negatives). The Time to Respond to Threats has been measured to 719 ms (average) from the time anomaly is detected to the time the Mitigator initiates the mitigation. A demonstration video of the results is made available [7].

V. CONCLUSION

We successfully demonstrated the role of decentralized anomaly detection in Kubernetes computing clusters to minimize disruption of 5G service delivery in case of cyber-attacks. Two types of DoS attacks have been selected to showcase the potential impact on 5G connectivity, and on the importance to detect misbehavior not only on external cluster communication but also on internal per worker and per pod level granularity. Anomaly detection on such level of traffic granularity provides higher level of assurance on the healthy status of clusters. Future work will focus on extending the PoC to 5G/6G service provisioning across multi-cluster and multi-slice environments on attack detection and mitigation.

ACKNOWLEDGMENT

This work received funding from the National Spanish MINECO under grant No. TSI-063000-2021-55 (6GDawn-RESILIENT), grant No. PID2021-126431OB-I00 funded by MCIN/AEI/ 10.13039/501100011033 (ANEMONE), the European Union's Horizon Europe research and innovation program under Grant Agreements No.101095542 (CYLCOMED) and No.101070537 (CROSSCON), and the EU's Digital Europe Program under GA No. 101190370 (SAFE).

REFERENCES

- [1] A. Chtioui, "The Evolution of Intrusion Detection Systems: Embracing Kubernetes and AI for Modern Security," Tisalabs 2024, Available at <https://www.tisalabs.com/2024/10/09/the-evolution-of-intrusion-detection-systems-embracing-kubernetes-and-ai-for-modern-security>
- [2] A. K. Bhardwaj, P. Dutta, & P. Chintale, "AI-Powered Anomaly Detection for Kubernetes Security: A Systematic Approach to Identifying Threats". *Babylonian Journal of Machine Learning*, 2024.
- [3] J. Kosińska and M. Tobiasz, "Detection of Cluster Anomalies With ML Techniques," in *IEEE Access*, vol. 10, pp. 110742-110753, 2022.
- [4] J. G. Almaraz-Rivera, "An Anomaly-based Detection System for Monitoring Kubernetes Infrastructures," in *IEEE Latin America Transactions*, vol. 21, no. 3, pp. 457-465, March 2023.
- [5] <https://www.srsran.com/5g>
- [6] <https://booklet.evidenresearch.eu/lads>
- [7] <https://booklet.evidenresearch.eu/lads#:~:text=Demonstration%20videos> or <https://booklet.evidenresearch.eu/sites/booklet/files/public/content-video/2025/6gdawn-resilient-uc1-poc1-decentralised-intrusion-detection-final-demonstration-v3s.mp4>

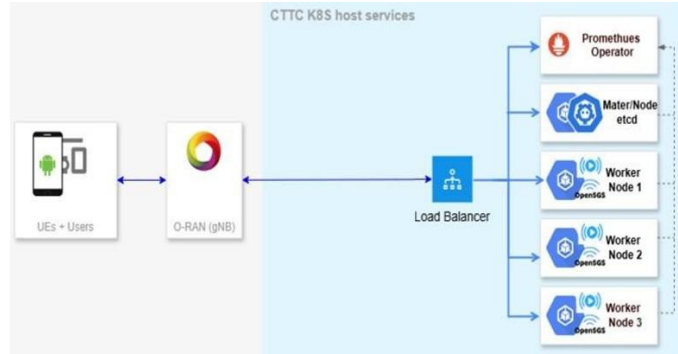


Figure 1. 5G Testbed Architecture

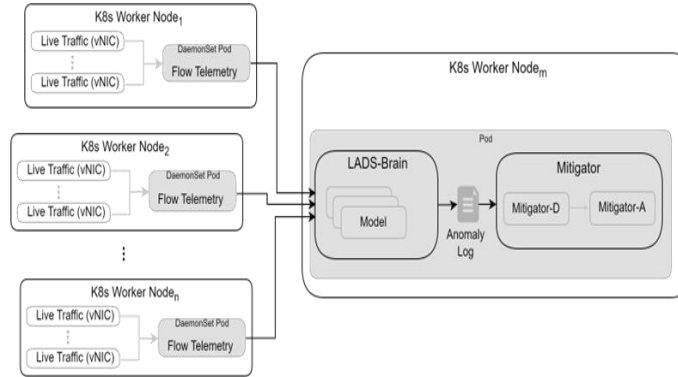


Figure 2. Intrusion detection and mitigation workflow in a Kubernetes cluster

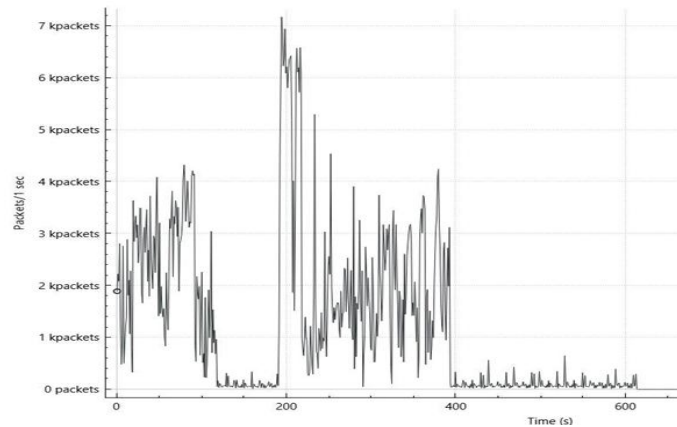


Figure 3. Traffic footprint of the blackholing through ARP poisoning on the UPF service

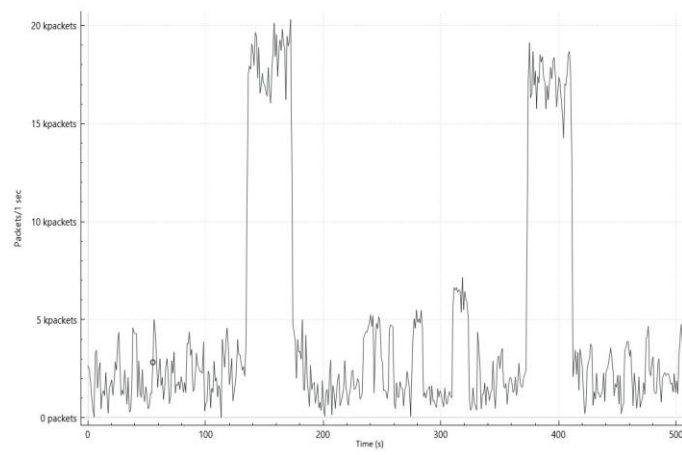


Figure 4. Traffic footprint of the ICMP ping flood on the UPF