UWS Academic Portal

Hardware-accelerated firewall for 5G mobile networks

Ricart-Sanchez, Ruben; Malagon, Pedro; Alcaraz-Calero, Jose M.; Wang, Qi

Published in:
2018 IEEE 26th International Conference on Network Protocols (ICNP)

DOI:
[10.1109/ICNP.2018.00066](https://doi.org/10.1109/ICNP.2018.00066)

Published: 25/09/2018

Document Version
Peer reviewed version

[Link to publication on the UWS Academic Portal](#)

Citation for published version (APA):
Ricart-Sanchez, R., Malagon, P., Alcaraz-Calero, J. M., & Wang, Q. (2018). Hardware-accelerated firewall for 5G mobile networks. In *2018 IEEE 26th International Conference on Network Protocols (ICNP)* (pp. 446-447). (IEEE Conference Proceedings). IEEE. <https://doi.org/10.1109/ICNP.2018.00066>

General rights

Copyright and moral rights for the publications made accessible in the UWS Academic Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact pure@uws.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

“© © 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Hardware-Accelerated Firewall for 5G Mobile Networks

1st Ruben Ricart-Sanchez

*School of Computer and Engineering
University of the West of Scotland
Paisley, United Kingdom
Ruben.Ricart-Sanchez@uws.ac.uk*

2nd Pedro Malagon

*Departamento de Ingenieria Electronica
Universidad Politecnica de Madrid
Madrid, Spain
pedro.malagon.marzo@upm.es*

3rd Jose M. Alcaraz-Calero

*School of Computer and Engineering
University of the West of Scotland
Paisley, United Kingdom
Jose.Alcaraz-Calero@uws.ac.uk*

4th Qi Wang

*School of Computer and Engineering
University of the West of Scotland
Paisley, United Kingdom
Qi.Wang@uws.ac.uk*

Abstract—The evolution from the current Fourth-Generation (4G) networks to the emerging Fifth-Generation (5G) technologies implies significant changes in the architecture and poses demanding requirements on network infrastructures. One of the Key Performance Indicators (KPIs) in 5G is to ensure a secure network with zero downtime. In this paper, we focus on the provisioning of protection capabilities for 5G infrastructures. Our objective is to implement a new 5G firewall that allows the detection, differentiation and selective blocking of 5G network traffic in the edge-to-core network segment of a 5G infrastructure, using a hardware-accelerated framework based on Field Programmable Gate Arrays (FPGA), developed using the P4 language. The proposed 5G firewall has been prototyped with the new capabilities proposed empirically validated.

Index Terms—P4-NetFPGA, FPGA, 5G mobile networks, network protection, firewall

I. INTRODUCTION

Mobile telecommunication operators should be ready to manage a high volume of network traffic in 5G networks. This entails the development of new 5G network elements that are able to manage those volumes of data. This paper is focused on the development of a new 5G firewall. The architecture proposed for this solution is based on a 5G Mobile Edge Computing (MEC) architecture where the Radio Access Network (RAN) is deployed on an Edge location to allow fast processing of data in the last mile, close to the final user and where the Edge is connected to the Core network segment to get access to other users globally. The proposed 5G firewall is located between the Edge and the Core network segment to be able to drop malicious traffic of 5G mobile users in order to protect both 5G users and infrastructure.

To the best of our knowledge, there is no similar existing solution providing such capabilities. To demonstrate our contribution, a full prototype has been implemented and extensive empirical validation of the prototype has been carried out using a 5G Edge-to-Core infrastructure currently deployed in our data centre. The main output of this contribution is a prototype of a new hardware-accelerated 5G firewall for the new 5G

traffic, based on a NetFPGA Network Interface Card (NIC) [1] that makes use of the P4 language [2].

II. PROPOSED P4-NETFPGA-BASED 5G FIREWALL ARCHITECTURE

Fig. 1. 5G architecture overview

Figure 1 shows the 5G Edge-Core architecture deployed in this research work. The main functionality of the 5G firewall is to protect the internal 5G infrastructure of a communication operator. The 5G infrastructure requires the network traffic over it to fulfill the requirements of traffic encapsulation to allow user mobility across all the elements of the infrastructure, mainly by providing support for the General Packet Radio Service (GPRS) Tunneling Protocol (GTP) in a Long-Term Evolution (LTE) based 5G network. The corresponding 5G packets have the following header structure: MAC, IPv4 (outer), UDP/TCP (outer), GTP, IPv4 (inner), and UDP/TCP (inner).

It is noted that conventional firewalls are mainly focused on the classification of traditional IP traffic and thus only inspect the outer IP and UDP/TCP headers. In contrast, the proposed solution goes further in the traffic classification and is able to

inspect the inner IP and UDP/TCP headers, which are those related to the 5G users in the Edge-to-Core network segment.

The proposed FPGA-based 5G firewall adopts the three stages of the P4-NetFPGA pipeline: parser, match/action and deparser, as explained below.

1) *Parsing stage*: Five different headers are defined in order to allow their parsing inside of the P4 program: MAC, IP, UDP/TCP and GTP.

2) *Match/Action pipeline*: After the parsing stage, if the packet has the structure expected by the parser, it enters the match/action pipeline, where the data extracted from the packet are used to decide if the packet should be dropped or not. In this solution, a Ternary Content Addressable Memory (TCAM) has been implemented in P4. It is composed of several keys: 5G User Source IP, 5G User Destination IP, 5G User Source Port, 5G User Destination Port, Transport Protocol Type, and the identification number of the GTP tunnel. This TCAM table allows storing the 5G firewall rules, which determines if an action should be applied. The prototype allows the DROP action to be enforced to rule-defined malicious packets and has an allow-by-default policy.

3) *Deparser*: At this stage, if the packet has not be dropped, it is rebuilt and transmitted through the 5G infrastructure.

III. EXPERIMENTAL SETUP

A realistic use case has been deployed to empirically validate the prototyped P4-NetFPGA-based 5G firewall. The use case implies the prototyping of a user-level Application Programming Interface (API) to allow a network administrator to populate firewall rules from the user space into the NetFPGA-SUME board. This API uses the existing SDK libraries provided by the NetFPGA-SUME project [3], allowing us to insert or remove rules from the NetFPGA.

Fig. 2. Experimental testbed to validate the prototyped 5G firewall.

Figure 2 shows the experimental testbed developed to validate the 5G firewall prototype. Rules are inserted when some 5G flows should be blocked. These rules can be inserted at run-time when the NetFPGA is receiving the traffic, although for this experimental setup, the rules are inserted before the transmission starts (steps 1 and 2 in Figure 2). For the execution of the experiments, 10 different pcaps have been created. They contain real traffic of up to 512 different 5G users where each of the users is generating two different flows.

The idea is to validate the prototyped 5G firewall when up to 1024 flows are arriving at the same time.

As shown in Figure 2, the traffic sent starts in PC1 (Edge) and it goes to PC2 (Core). The non-blocked traffic will come back to PC1. This return is only for experimental setup because it allows the measurement of the delay applied by the prototype in the NetFPGA card. Pcaps are sent separately, one by one, from PC1 (step 3 in Figure 2). The packets of each pcap go through the eth1 interface to the nf0 interface in PC2 (step 4 in Figure 2). When packets received are not discarded in the NetFPGA, they are forwarded through the nf1 to PC1 (step 5 in Figure 2). Finally, with the packets sent and the packets received, both captured at PC1, a comparison is conducted to empirically validate this solution (step 6 in Figure 2). The testing scenario is that one out of the two flows of each of the users is be dropped while the other one is be allowed.

In all the validation tests conducted, the 5G firewall prototype was able to deal with up to simultaneous 1024 flows, and block the flows that were identified as malicious by the inserted rules correctly, showing a high degree of reliability of the proposed solution.

IV. CONCLUSION

This paper has provided a new hardware-accelerated 5G firewalls suitable for deployment in the Edge-to-Core 5G network segment. The proposed solution is able to protect 5G infrastructures and mitigate cyber-attacks by selectively blocking malicious traffic that has been identified, due to the fact that it is implemented on the border of the core network and it receives all traffic of the mobile operator. Experimental results have validated the design and prototyping of the proposed 5G firewall system. Furthermore, the prototype yields good reliability even in stressed scenarios.

Future work will further investigate and quantify the performance of the proposed solution.

ACKNOWLEDGMENT

This work was funded by the European Commission Horizon 2020 5G-PPP Program under Grant Agreement Number H2020-ICT-2016-2/761913 (SliceNet: End-to-End Cognitive Network Slicing and Slice Management Frame-790 work in Virtualised Multi-Domain, Multi-Tenant 5G Networks), and by the Spanish Ministry of Economy and Competitiveness under contract RTC-2016-5434-8.

REFERENCES

- [1] J. W. Lockwood, N. McKeown, G. Watson, G. Gibb, P. Hartke, J. Naous, R. Raguraman, and J. Luo, "Netfpga—an open platform for gigabit-rate network switching and routing," in *Microelectronic Systems Education, 2007. MSE'07. IEEE International Conference on*. IEEE, 2007, pp. 160–161.
- [2] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker, "P4: Programming protocol-independent packet processors," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 87–95, Jul. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2656877.2656890>
- [3] N. Zilberman, Y. Audzevich, G. A. Covington, and A. W. Moore, "Netfpga sume: Toward 100 gbps as research commodity," *IEEE micro*, vol. 34, no. 5, pp. 32–41, 2014.