Newsletter

SMARTNESS

No. 5 – Dec/24

Table of Contents

- Remarkable participation at NFV-SDN 2024, Natal, Brazil
- New visit of Prof. Fredrik Heintz from Linköping University, Sweden
- Participation at 6G Briefing International Conference
- PhD Symposium Award at IEEE NFV-SDN 2024
- Participation at SBSeg 2024, São José dos Campos, Brazil
- SMARTNESS 2030 Research Center and Open Science
- WARA-PS event and Technical Visit to Sweden
- SMARTNESS 2030 and 6G at Futurecom
- Short course on AI for Networking at UdelaR, Uruguay
- New publication presented at CloudNet 2024

Keep up to date with our latest publications, events, news and more!

Follow us on our social media:

smartness2030.tech

smartness2030

smartness2030.tech

smartness2030

Remarkable participation at NFV-SDN 2024, Natal, Brazil

SMARTNESS participated in the IEEE NFV-SDN 2024 conference in Natal, Brazil, presenting cutting-edge research on Network Function Virtualization (NFV) and Software-Defined Networking (SDN). The team showcased two papers, three demos, a doctoral symposium, a tutorial, and a keynote. Highlights include the award-winning PINT-BoX demo, RESISTING for resilient networking, MATADOR for cloud gaming traffic detection, and key insights on 6G research and programmable infrastructures.

Read more

New visit of Prof. Fredrik Heintz from Linköping University

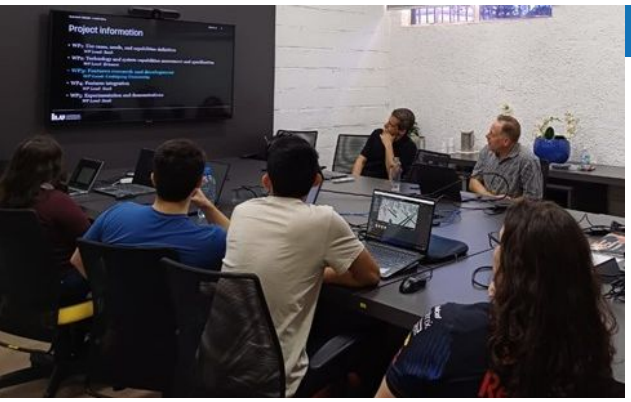

On November 14, Prof. Luiz Bittencourt hosted Prof. Fredrik Heintz and PhD student Amath Sow from Linköping University at UNICAMP's Institute of Computing. They joined Ericsson Research Brazil, UFABC, and UNICAMP researchers to discuss future collaborations on drone collision avoidance, video streaming, and vehicle simulation. Key research updates were shared, fostering plans for joint papers and international conference submissions.

Read more

Participation at 6G Briefing International Conference

Marina Martinelli, PhD researcher at the Institute of Geosciences and SMARTNESS 2030, attended the 6G Briefing 2024 event, connecting academia and industry experts. The conference featured talks on 6G, XR, V2X, Ubiquity, and Quantum Computing, promoting R&D collaboration and innovative strategies. Diversity, inclusion, and cutting-edge tech were key highlights, fostering competitive growth in research and industry.

Read more

PhD Symposium Award at IEEE NFV-SDN 2024

PhD candidate Md Tariqul Islam won the Best Doctoral Symposium Paper award for his work “QoE Evaluation for Emerging Media Applications: Network-Level Analysis and Traffic Modeling.” His research tackles ensuring Quality of Experience (QoE) in VR/AR by analyzing latency, bandwidth, and stability. His model helps network operators optimize traffic for immersive experiences, advancing QoE research in 5G and beyond.

Read more

Participation at SBSeg 2024, São José dos Campos, Brazil

The SBSeg 2024, Brazil's top forum on information and computer systems security, was held in São José dos Campos from September 16-19. SMARTNESS researchers stood out with notable contributions. PhD student Caio Teixeira (UNICAMP) won Best Short Article in the cryptography track for his work on post-quantum encryption in vehicular networks. M.Sc. student Rodrigo Pierini, alongside Professors Rothenberg and Henriques, presented research on implementing the Forro14 cypher on Tofino switches using P4. USP student Thiago Ferreira earned an honourable mention for applying fuzzing in SPDM testing, co-authored by Professor Simplicio Jr. Additionally, Professor Batista was honoured for his reviewer contributions.

Read more

SMARTNESS 2030 Research Center and Open Science

The SMARTNESS 2030 Center is committed to research excellence and the broad dissemination of scientific knowledge, supporting Open Science principles defined by FAPESP. It ensures public access to research results while safeguarding privacy, security, and intellectual property. Its Open Science actions include branding, communication, events, training, and publications. Expert talks, like those by Dr. Claudia Bauzer Medeiros and Dr. Eduardo Sartori, strengthen its practices in research ethics, data management, and Open Data dissemination.

Read more

WARA-PS event and Technical Visit to Sweden

Professor Luiz Bittencourt, a SMARTNESS researcher, received funding from CISB to participate in the WARA-PS Demonstration Week 2024 in Sweden. He presented a demo on drone management using edge computing, focusing on algorithmic strategies for collision avoidance and take-off and landing management. He also made technical visits to Professor Fredrik Heintz's lab at Linköping University, Saab's headquarters in Linköping, and Ericsson Research in Stockholm.

Read more

SMARTNESS 2030 and 6G visions at Futurecom

Marina Martinelli, PhD researcher at SMARTNESS 2030 (UNICAMP), spoke at Futurecom about 6G's transformative potential beyond faster internet, highlighting its role in telemedicine, autonomous vehicles, and immersive technologies like AR and VR. She stressed 6G's power to bridge digital divides using low-orbit satellites and relay drones, enhancing services in remote areas. The panel also addressed smart cities, sustainability, quantum networks, and ITU's 2024 standardization efforts.

Read more

Short course on AI for Networking at UdelaR, Uruguay

Read more

Prof. Rafael Pasquini and Prof. Christian Rothenberg contributed to the Workshop on Artificial Intelligence in Cyber-Physical Systems, held at the Faculty of Engineering at the Universidad de la República (UdelaR) in Uruguay. The event was organized to encourage the intersection of AI, machine learning (ML), and networking technologies. Over two intensive days, the short course on AI for Networking workshop covered 12 hours of theoretical knowledge, hands-on activities, and practical examples.

New publication presented at CloudNet 2024

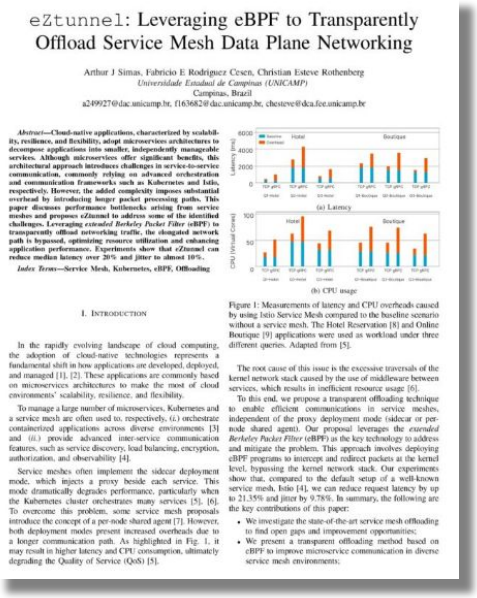

eZtunnel: Leveraging eBPF to Transparently Offload Service Mesh Data Plane Networking

Arthur J Simas, Fabricio E Rodriguez Cesen, Christian Esteve Rothenberg
Universidade Estadual de Campinas (UNICAMP)
Campinas, Brazil
a249927@dac.unicamp.br, f163682@dac.unicamp.br, chesteve@dca.fee.unicamp.br

Abstract—Cloud-native applications, characterized by scalability, resilience, and flexibility, adopt microservices architectures to decompose applications into smaller, independently manageable services. Although microservices offer significant benefits, this architectural approach introduces challenges in service-to-service communication, commonly relying on advanced orchestration and communication frameworks such as Kubernetes and Istio, respectively. However, the added complexity imposes substantial overhead by introducing longer packet processing paths. This paper discusses performance bottlenecks arising from service meshes and proposes eZtunnel to address some of the identified challenges. Leveraging *extended Berkeley Packet Filter* (eBPF) to transparently offload networking traffic, the elongated network path is bypassed, optimizing resource utilization and enhancing application performance. Experiments show that eZtunnel can reduce median latency over 20% and jitter to almost 10%.

Index Terms—Service Mesh, Kubernetes, eBPF, Offloading

I. Introduction

In the rapidly evolving landscape of cloud computing, the adoption of cloud-native technologies represents a fundamental shift in how applications are developed, deployed, and managed [1], [2]. These applications are commonly based on microservices architectures to make the most of cloud environments' scalability, resilience, and flexibility.

To manage a large number of microservices, Kubernetes and a service mesh are often used to, respectively, (*i.*) orchestrate containerized applications across diverse environments [3] and (*ii.*) provide advanced inter-service communication features, such as service discovery, load balancing, encryption, authorization, and observability [4].

Service meshes often implement the sidecar deployment mode, which injects a proxy beside each service. This mode dramatically degrades performance, particularly when the Kubernetes cluster orchestrates many services [5], [6]. To overcome this problem, some service mesh proposals introduce the concept of a per-node shared agent [7]. However, both deployment modes present increased overheads due to a longer communication path. As highlighted in Fig. 1, it may result in higher latency and CPU consumption, ultimately degrading the Quality of Service (QoS) [5].

Figure 1: Measurements of latency and CPU overheads caused by using Istio Service Mesh compared to the baseline scenario without a service mesh. The Hotel Reservation [8] and Online Boutique [9] applications were used as workload under three different queries. Adapted from [5].

The root cause of this issue is the excessive traversals of the kernel network stack caused by the use of middleware between services, which results in inefficient resource usage [6].

To this end, we propose a transparent offloading technique to enable efficient communications in service meshes, independent of the proxy deployment mode (sidecar or per-node shared agent). Our proposal leverages the *extended Berkeley Packet Filter* (eBPF) as the key technology to address and mitigate the problem. This approach involves deploying eBPF programs to intercept and redirect packets at the kernel level, bypassing the kernel network stack. Our experiments show that, compared to the default setup of a well-known service mesh, Istio [4], we can reduce request latency by up to 21.35% and jitter by 9.78%. In summary, the following are the key contributions of this paper:

- We investigate the state-of-the-art service mesh offloading to find open gaps and improvement opportunities;
- We present a transparent offloading method based on eBPF to improve microservice communication in diverse service mesh environments;

Read more

Christian Esteve Rothenberg, Dr. Fabricio Rodríguez Cesen, and M.Sc. student Arthur J. Simas (UNICAMP) had their paper “eZtunnel: Leveraging eBPF to Transparently Offload Service Mesh Data Plane Networking” accepted at IEEE CloudNet 2024 in Rio de Janeiro. The study addresses service mesh bottlenecks using eBPF to offload traffic, reducing network path length, improving resource use, and cutting median latency by over 20% and jitter by ~10%.

ACKNOWLEDGMENTS & DISCLAIMER

This work has been performed within the framework of the FAPESP Engineering Research Center (ERC) Program under FAPESP grant agreement #2021/00199-8 (SMARTNESS).

The information in this document reflects the SMARTNESS ERC's view, but the partner institutions of SMARTNESS are not liable for any use that may be made of any of the information contained therein. The views and opinions expressed are those of the author(s) only and do not necessarily represent those of FAPESP or the other granting authorities. Neither FAPESP nor the granting authority can be held responsible for them. The views expressed are solely those of the authors and do not necessarily represent Ericsson's official standpoint.

Keep up to date with our latest publications, events, news and more!

Follow us on our social media:

smartness2030.tech

smartness2030

smartness2030.tech

smartness2030