# Lightweight low-latency virtual networking

Final talk for the Bachelor's Thesis by

**Alexander Daichendt**
advised by Florian Wiedner, Jonas Andre

Thursday 8th September, 2022

Chair of Network Architectures and Services
Department of Informatics
Technical University of Munich

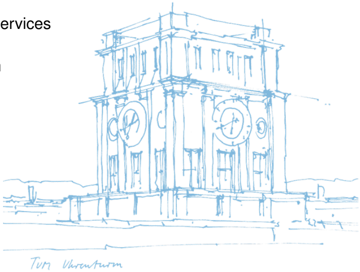

# Research Questions

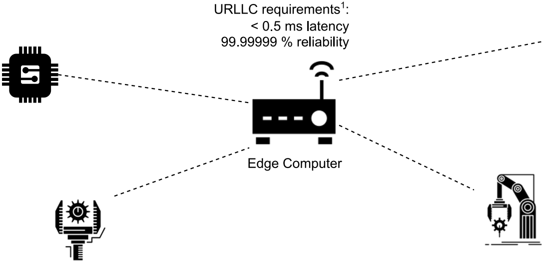


- **RQ1**: Are containers viable for URLLC[1]?
- **RQ2**: How large is the difference in network latencies between containers and VMs?
- **RQ3**: What are the root causes for the differences, and what can we do to minimize them?

[1] S. Gallenmüller et al. "5G URLLC: A Case Study on Low-Latency Intrusion Prevention". In: IEEE Communications Magazine 58.10 (Okt. 2020), S. 35–41. DOI: 10.1109/MCOM.001.2000467.

# Background information

## Containers

- shares kernel with host
- no abstraction layers for system calls and paging → more performance
- rapid deployment and startup

⇒ lightweight virtualization

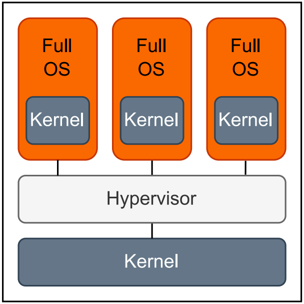


Virtual Machines

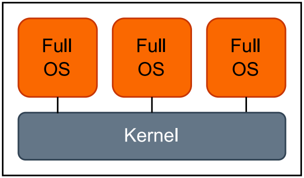


Containers

# Background information

## Relevant kernel-level features

**Control groups (cgroups)**: manage hardware resources
Examples:

- cpuset.cpus = 8-10 → limit cores
- memory.max = 34359738368 → memory limit in bytes
- pids.max = 1000 → limits creation of new processes

**Namespaces**: restricts a process' view on the system

**Performance Enhancement of Virtualized Media Gateway with DPDK for 5G Multimedia Communications**[2]

- network latencies with Docker and DPDK based media gateway for the 5G core network
- not directly comparable → adjusts Layer 3 and 4 information; no optimizations

**Ducked Tails: Trimming the Tail Latency of(f) Packet Processing Systems**[3]

- presents low-latency kernel optimizations → fundamental for our measurements
- compares latencies of a DPDK forwarder for bare metal and VM
- not directly comparable → lower packet rates

**HVNet: Hardware-Assisted Virtual Networking on a Single Physical Host**[4]

- virtual networking on a single host with VMs
- framework for our implementation

[2] W. Chen et al. "Performance Enhancement of Virtualized Media Gateway with DPDK for 5G Multimedia Communications". In: 2019 International Conference on Intelligent Computing and its Emerging Applications (ICEA). IEEE, Aug. 2019. DOI: 10.1109/icea.2019.8858303. URL: https://doi.org/10.1109/icea.2019.8858303.

[3] S. Gallenmüller et al. "Ducked Tails: Trimming the Tail Latency of(f) Packet Processing Systems". In: 3rd International Workshop on High-Precision, Predictable, and Low-Latency Networking (HiPNet 2021). Izmir, Turkey, Okt. 2021.

[4] F. Wiedner et al. "HVNet: Hardware-Assisted Virtual Networking on a Single Physical Host". In: IEEE INFOCOM WKSHPS: Computer and Networking Experimental Research using Testbeds (CNERT 2022) (INFOCOM WKSHPS CNERT 2022). Virtual Event, Mai 2022.

# Approach

## Integration of containers into pos - Setup

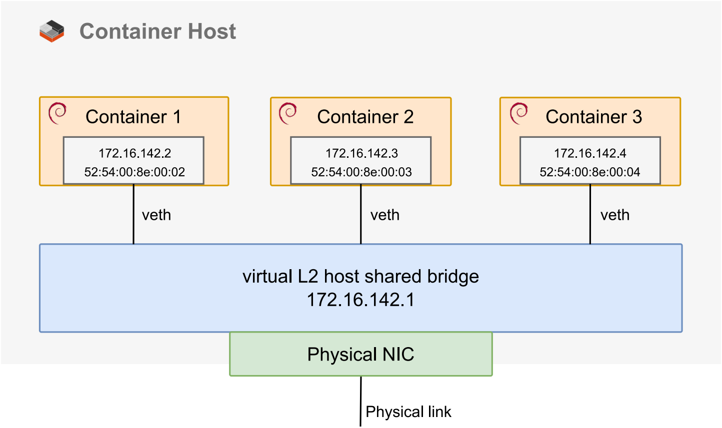

# Approach

## Integration of containers into pos - vBMC

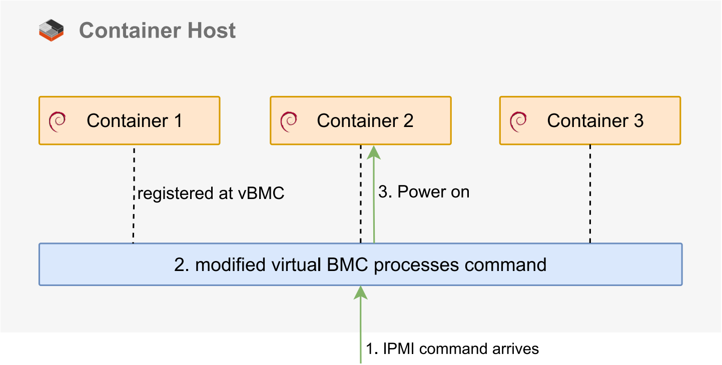

# Approach

## Integration of containers into HVNet

**Idea**:

- spawn containers instead of VMs
- execute the VM measurement scripts

**Challenges**:

- unable to load kernel modules in a container
- unable to create huge pages in a container

$\Rightarrow$ initialization of kernel-critical resources on container-host

# Measurements and results

## Experiment setup

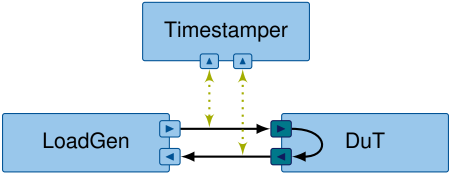


LoadGen:

- generates UDP traffic with MoonGen
- minimum packet size of 64 B

Timestamper:

- timestamps every single packet with passive optical TAPs
- writes result into pcap file

# Measurements and results

## Impact of interrupt recording (IR) on LXC

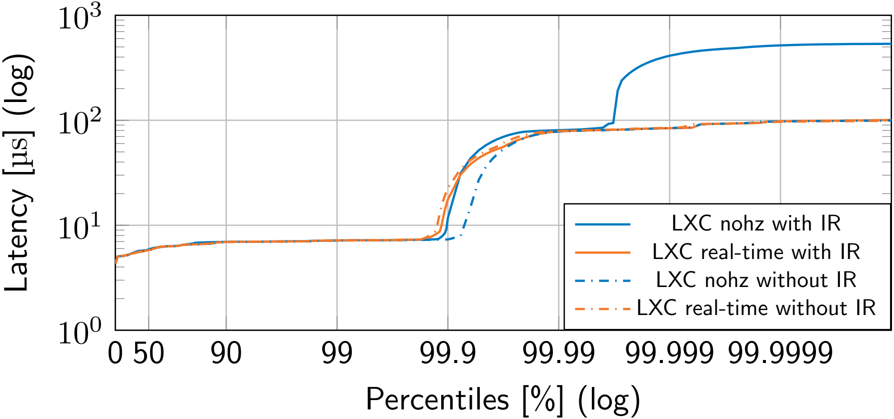

# Measurements and results

## base-nohz 1,52 Mpkt/s

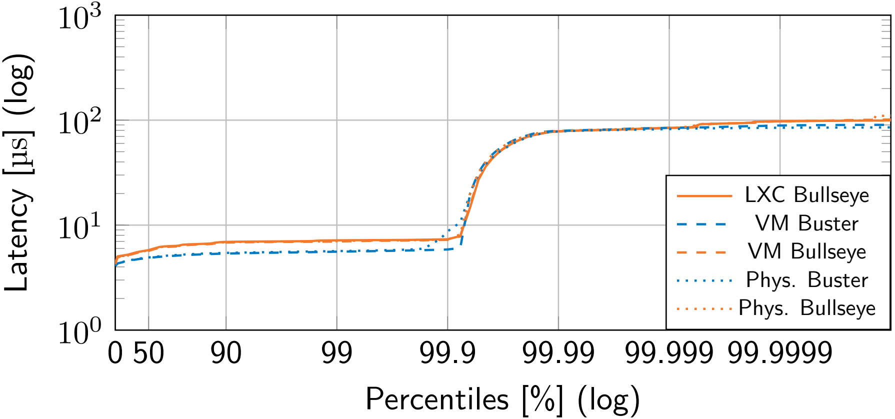

# Measurements and results

base-nohz 6,24 Mpkt/s

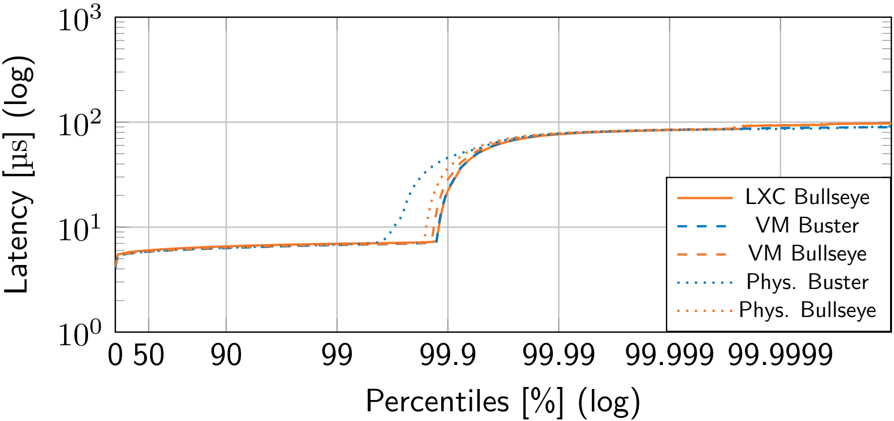

# Conclusion

## Summary

- integrated containers in testbed and implemented low-latency tooling
- LXC performs identical compared to VMs and bare metal
- average latencies of Debian Buster and Bullseye differ by 1,4 µs
- software running on the host influences LXC, but not VMs
- a real-time kernel stabilizes latencies

# Conclusion

## Comparison & Future work

Comparison:

| | LXC | MGW[5] | HiPNet[6] | HVNet[7] |
|---|---|---|---|---|
| VM | × | × | ✓ | ✓ |
| container | ✓ | ✓ | × | × |
| optimizations | ✓ | × | ✓ | ✓ |
| L3 + 4 adjustment | × | ✓ | × | × |
| Debian version | 11 | ○ | 10 | 10 |
| latency in µs; 99th perc. | 7.1 | > 69[8] | 3.3 | 5.5 |

Future work:

- CGroups v1 with Buster
- Experiments with SR-IOV
- Flow based experiments
- Comparison with other container solutions

5 W. Chen et al., “Performance Enhancement of Virtualized Media Gateway with DPDK for 5G Multimedia Communications”.

6 S. Gallenmüller et al., “Ducked Tails: Trimming the Tail Latency of(f) Packet Processing Systems”.

7 F. Wiedner et al., “HVNet: Hardware-Assisted Virtual Networking on a Single Physical Host”.

8 Paper specifies an average latency of 69 µs

# Bonus Slides

## Discussion preliminary results

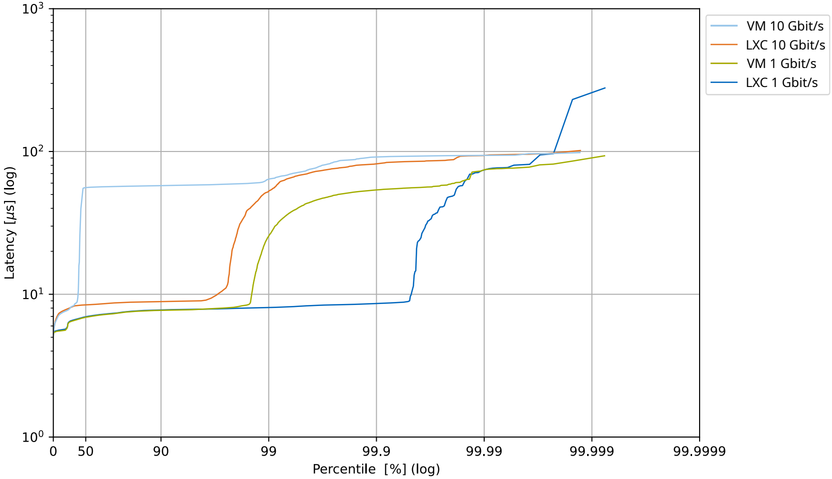

# Bonus Slides

## Low packet rates

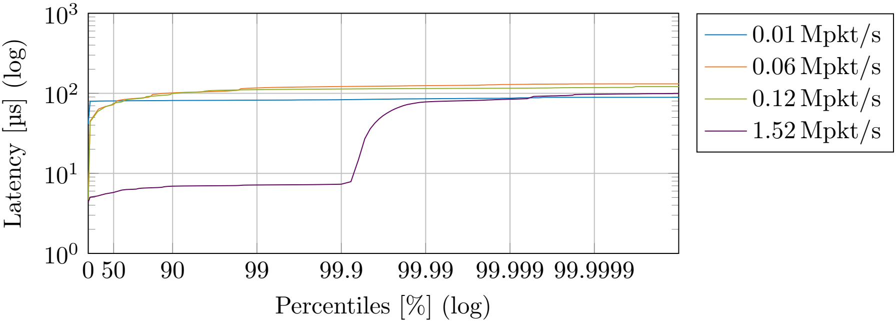

# Bonus Slides

## 5000 worst-case latencies for the base experiment with nohz kernel - 1,52 Mpkt/s

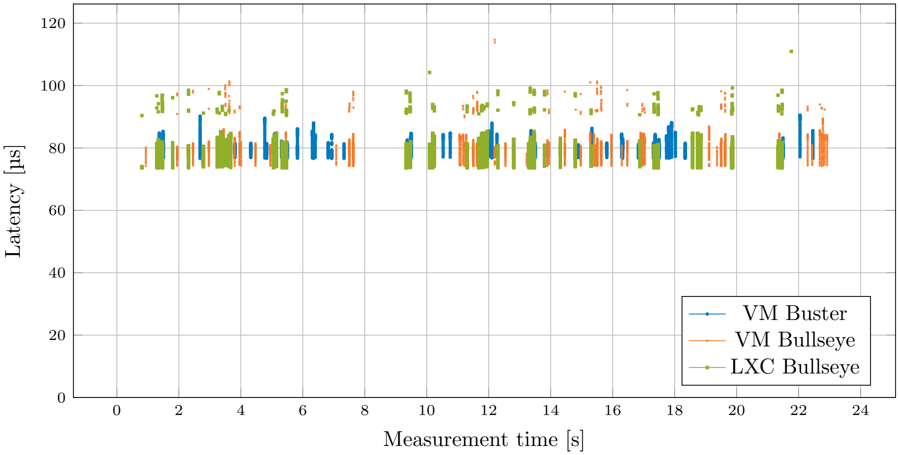

# Bonus Slides

## 5000 worst-case latencies for the base experiment with nohz kernel - 6,24 Mpkt/s