Virtual Open Systems

Cross-compartment Virtio-loopback: A bare-metal virtualization solution for the edge

Anna Panagopoulou – Alvise Rigo – Daniel Raho

ESARS2024 – Naples, November 2024

contact@virtualopensystems.com

www.virtualopensystems.com

NANCY

An Artificial Intelligent Aided Unified Network for Secure Beyond 5G Long Term Evolution

This work has received funding from the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101096456

Cross-compartment Virtio-loopback Scope

- Focus on Multi-access edge-computing (MEC) and Network Functions Virtualization (NFV)
- Cellular V2X scenarios to benefit from Beyond-SoTA NFV-based mechanisms
- Deployment of V2X services as VNFs at the network edge

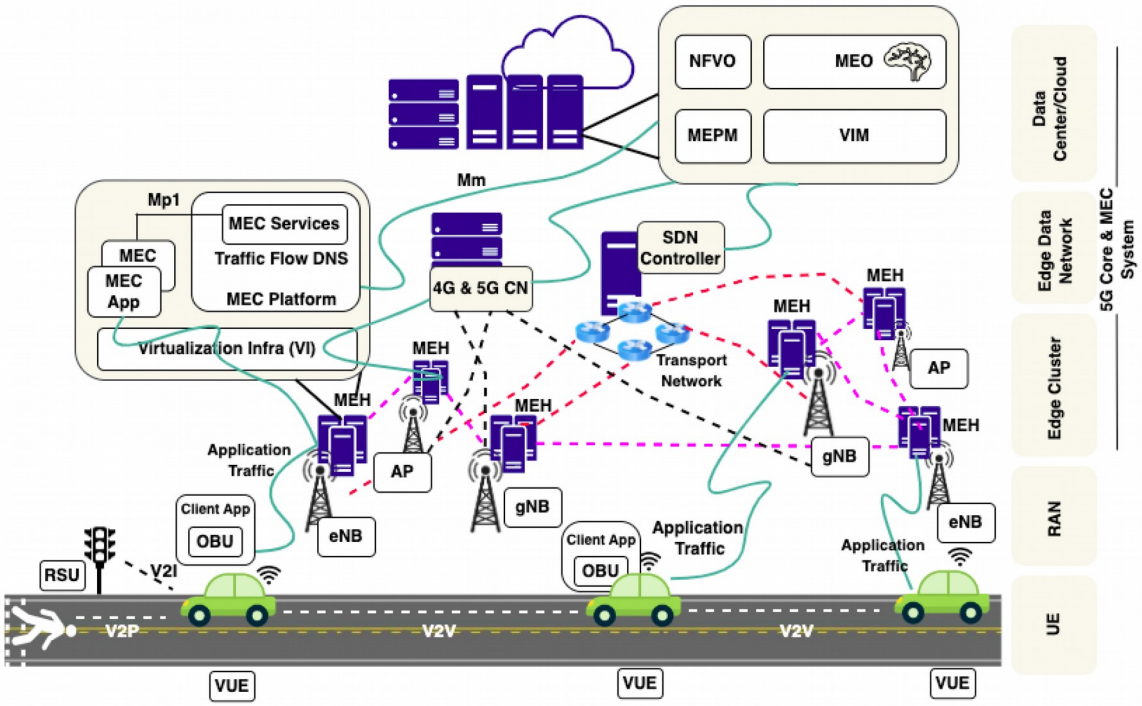

Cross-compartment Virtio-loopback Objectives

- Re-shape the virtualization environment at the edge
- Provide a **lightweight** and **hardware-isolated** virtualization solution for VNF deployments at edge servers
- Suitable for consolidating Ultra-reliable and Low-latency network services (ITS)

Cross-compartment Virtio-loopback Overview

- Lightweight, bare-metal virtualization solution, aimed for mixed-critical embedded architectures, such as the ARMv8 and the RISCv
- **Compartments**: Operating Systems with different critical levels, able to co-execute isolated in the same embedded hardware
- Reinforcement of edge devices with the ability to **host** the Virtualized Functions at individual Compartments, deployed in a **bare-metal fashion** on the physical hardware

Basic components- VOSySmonitor

- Highest-privilege low-level firmware for ARMv8 and RISCv
- Enables the concurrent execution of multiple Compartments
- Serves as the **partitioner** of the system resources in the mixed-critical world, and consolidates the different execution environments

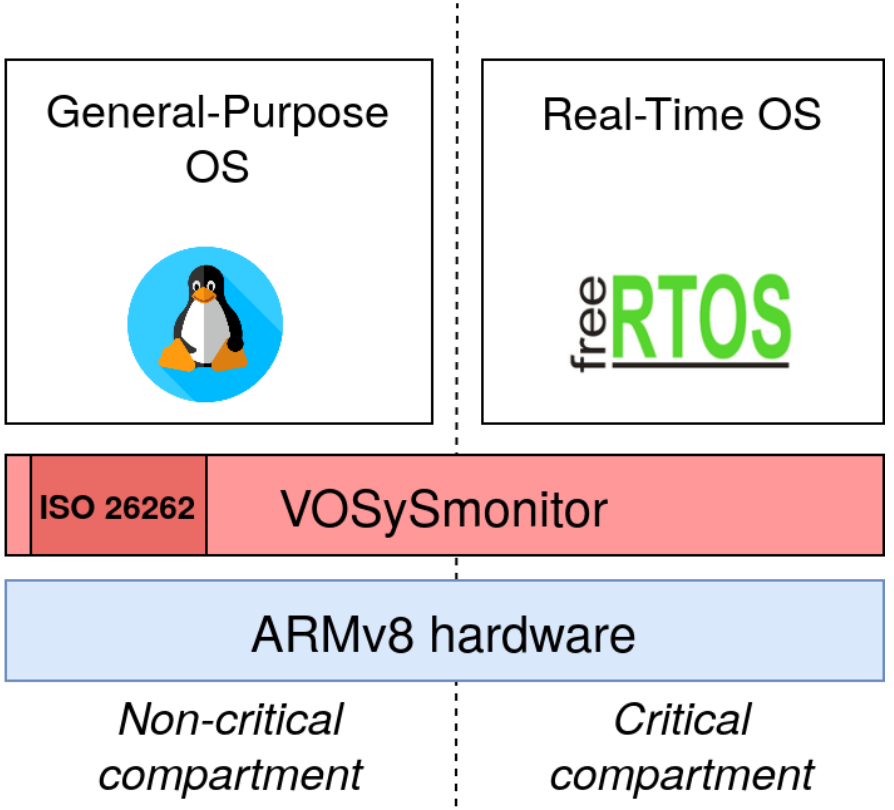

Basic components- VOSySVirtualnet

- Minimal **communication link** between Compartments in ARM
- Extends both VOSySmonitor as well as the OSes deployed on the Compartments
- Underlying mechanism to implement the generic Cross-Compartment communication

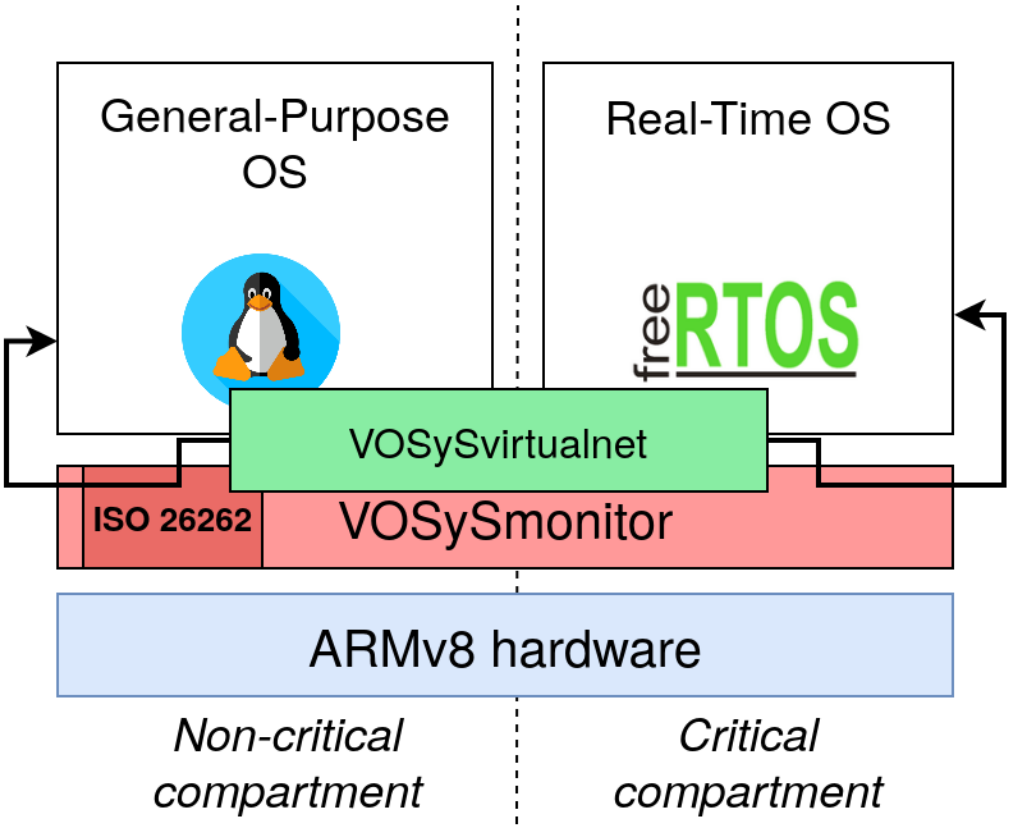

Basic components- Virtio Loopback

- Open-source HWL built for AGL, aiming to achieve re-usability of existing user-space applications into both Virtualized, Cloud and Bare-metal environments
- Introduces the Virtio standard into Bare-metal environments

- Builds connection between **Virtio drivers** on one end and **Vhost-user devices** on the other
- Sets the basis to decouple the Virtio driver from the Vhost-user device into separate systems

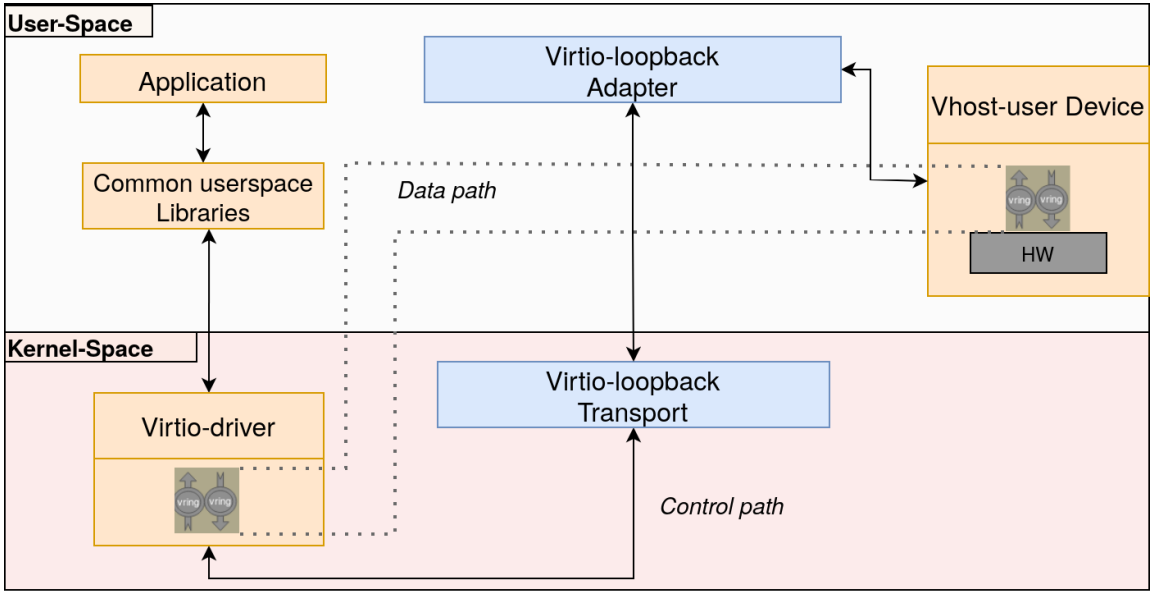

Cross-compartment Virtio-loopback Internals

- PoC with Critical and Non-critical Compartments
- Virtio-loopback introduction of **Client** and **Server** entities
- VOSySmonitor as the system **partitioner**
- VOSySvirtualnet for the **Control-path** of Virtio-loopback

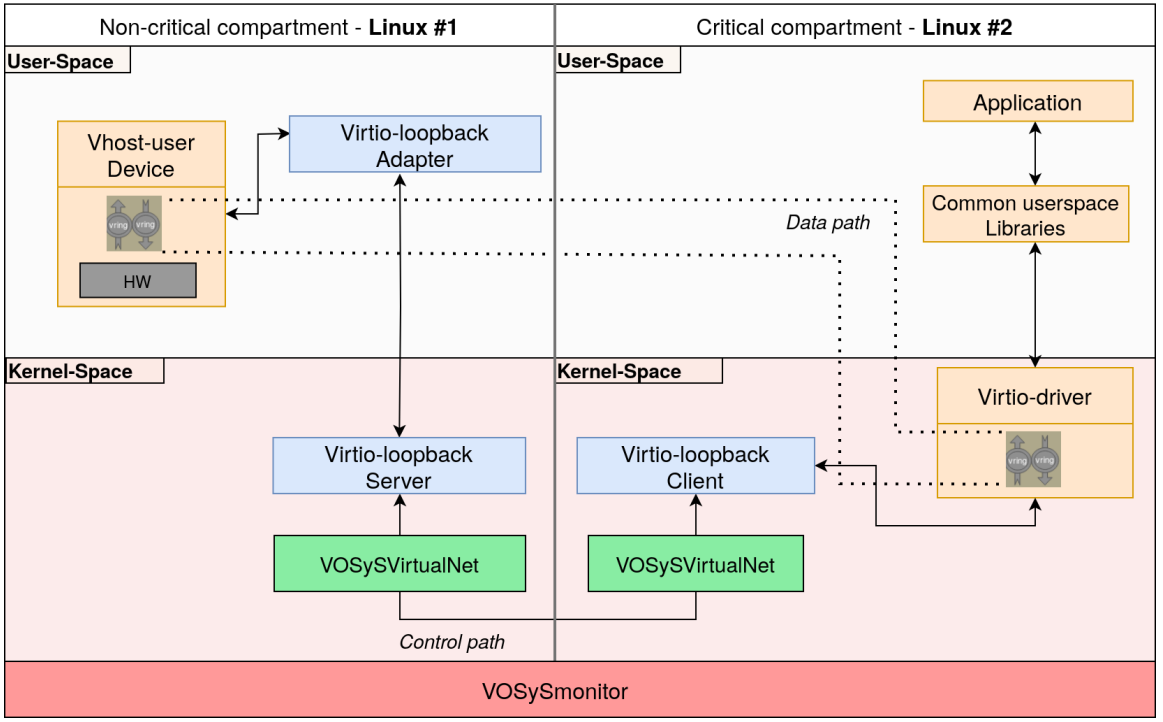

Experimentation environment

- Qemu virt VM setup
- Concurrent execution of two Linux images, one deployed in the non-Critical and the other in the Critical Compartment
- **Virtio-blk** device Fio benchmarks
- Comparing towards the original Virtio-loopback and also VDUSE

Fio benchmark results (1)

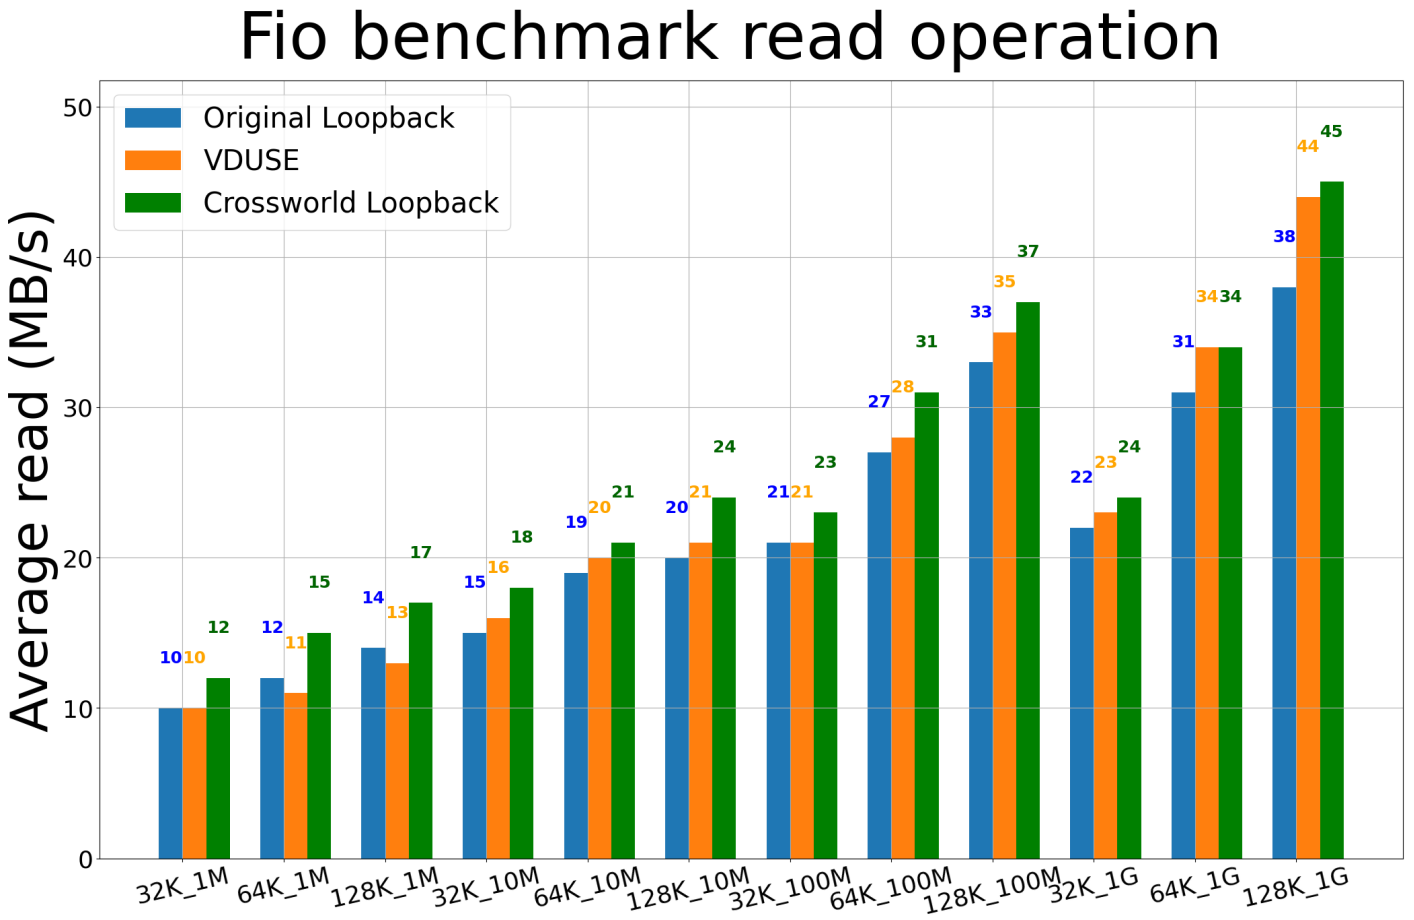

Fio benchmark results (2)

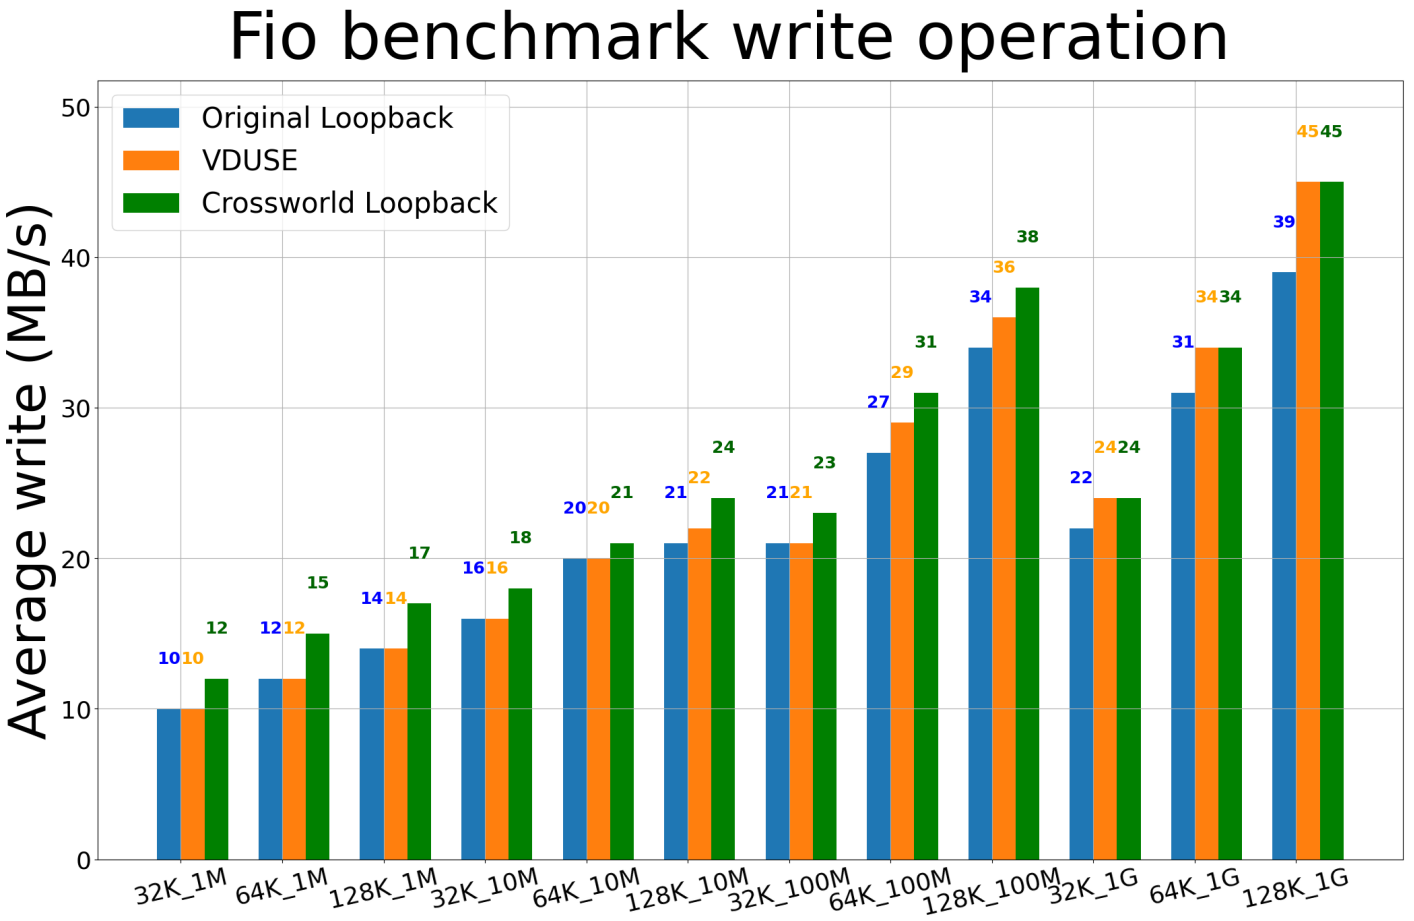

Conclusions (1)

- Promising results due to **decoupling** of the Client that initiates the requests, from the Server that asynchronously serves them
- **Increased data locality** in the data accesses
  - The data are consumed through a specified Shared-memory region between the Compartments

Conclusions (2)

- Proof that an abstracted VNF interfacing Virtio devices can be deployed to an individual ARMv8 compartment, in a **hardware-isolated** environment that is **spawned directly** in the bare-metal hardware.

- Consolidation of services with real-time and low latency demands, fitting well in the context of the modern,intelligent transportation services deployed in the networked infrastructure upon the cellular V2X interface.

Future directions

- Collection of benchmark results in an ARMv8 SoC
- Adjusting to RISCv architecture
- Integration with high-level network orchestrators (OpenStack and Compute Nova)

Thank you !

Questions ?

contact@virtualopensystems.com

Web: virtualopensystems.com

Products: http://www.virtualopensystems.com/en/products/

Demos: virtualopensystems.com/en/solutions/demos/

Guides: virtualopensystems.com/en/solutions/guides/

Research projects: virtualopensystems.com/en/research/innovation-projects/