Is Bare-metal I/O Performance with User-defined Storage Drives Inside VMs Possible?

Benchmarking libvfio-user vs. Common Storage Virtualization Configurations

Sebastian Rolon & Oana Balmau

CHEOPS '23
May 8th, 2023
Rome, Italy

Virtualization is everywhere

- Datasets keep growing
- We want storage to be efficient

Source: Google [2]. Inside a Google Datacenter.

Source: Intel [1]. An Intel Optane PCIe NVMe SSD.

aws

Azure

Google Cloud Platform

McGill UNIVERSITY

What is missing from storage virtualization?

1. Good performance with loose coupling
2. Rapid device prototyping
3. Live migration
4. Userspace drivers

McGill UNIVERSITY

What is missing from storage virtualization?

1. Good performance with loose coupling
2. Rapid device prototyping
3. Live migration
4. Userspace drivers

vfio-user

McGill UNIVERSITY

Can **vfio-user** be used as an alternative to current VM storage?

McGill UNIVERSITY

Virtualization and QEMU/KVM

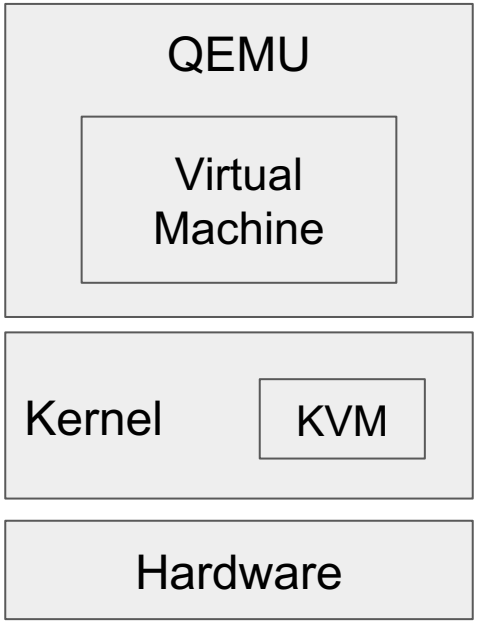

McGill UNIVERSITY

Context switches are expensive

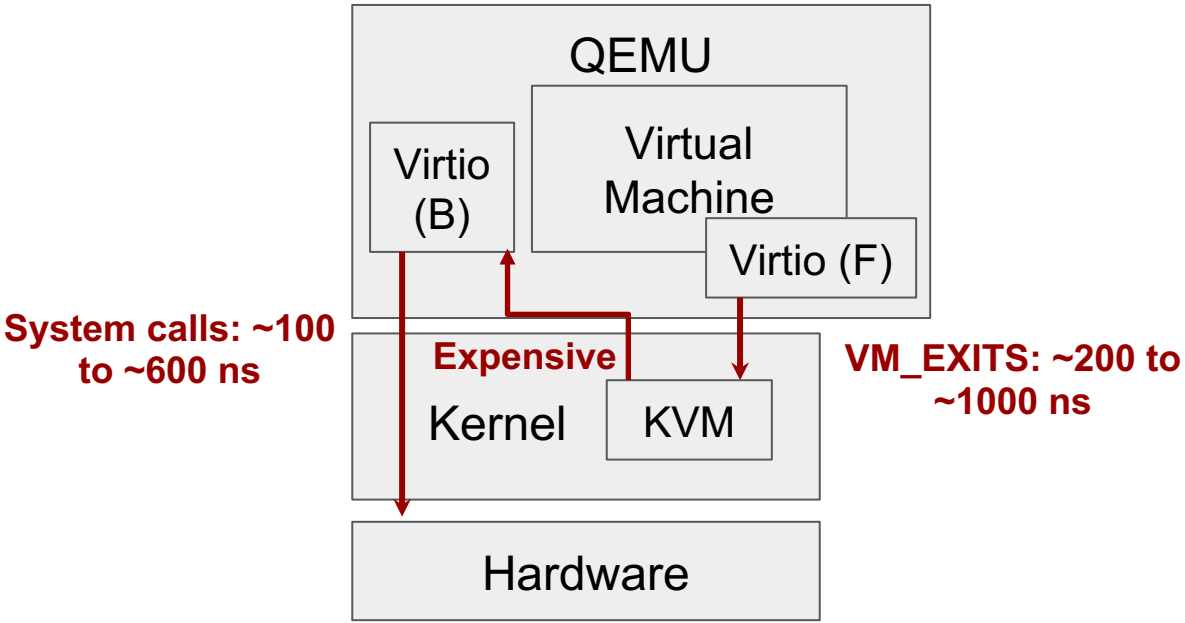

McGill UNIVERSITY

Userspace hardware access

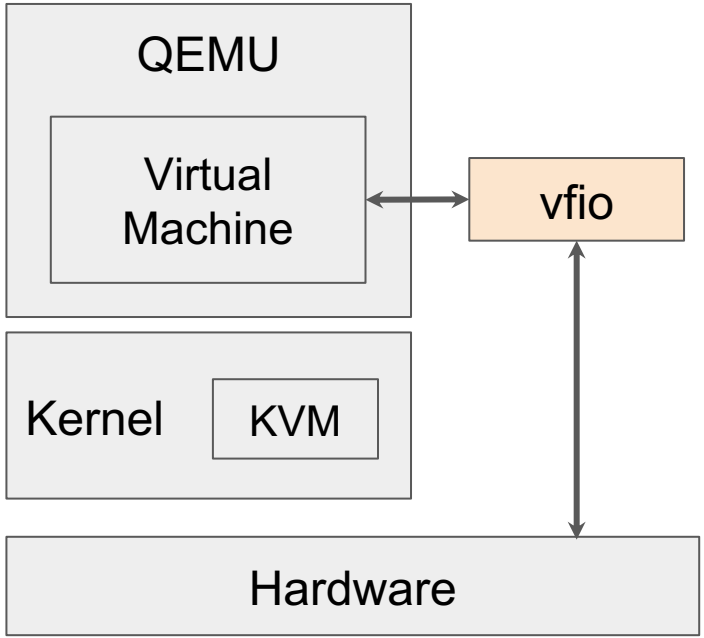

McGill UNIVERSITY

Abstracting hardware access

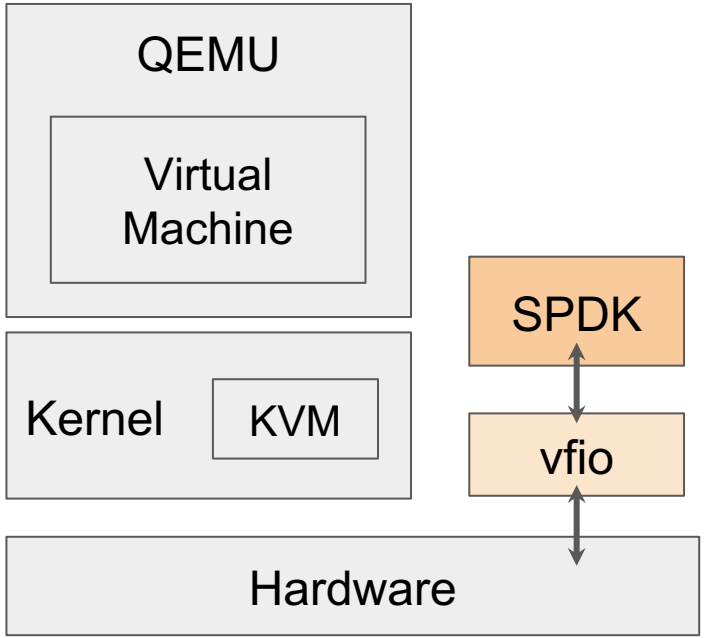

McGill UNIVERSITY

How do we connect after abstracting?

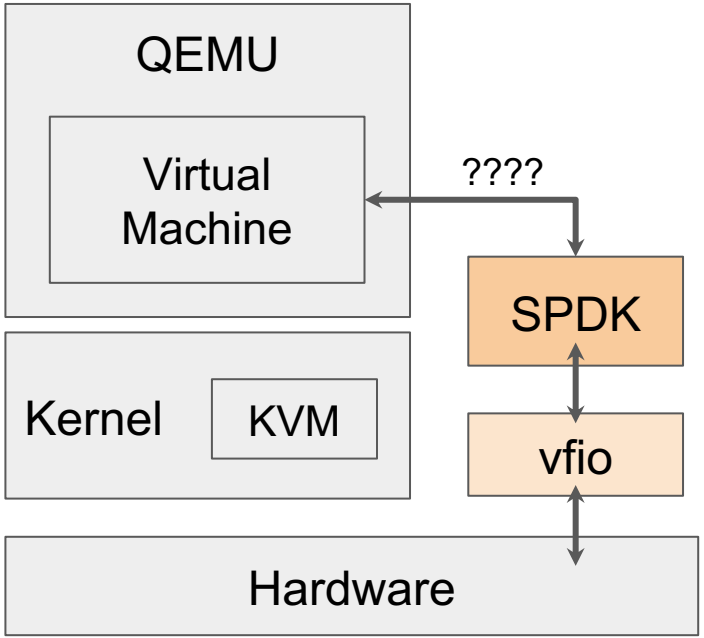

McGill UNIVERSITY

Vfio-user virtualizes hardware over a channel

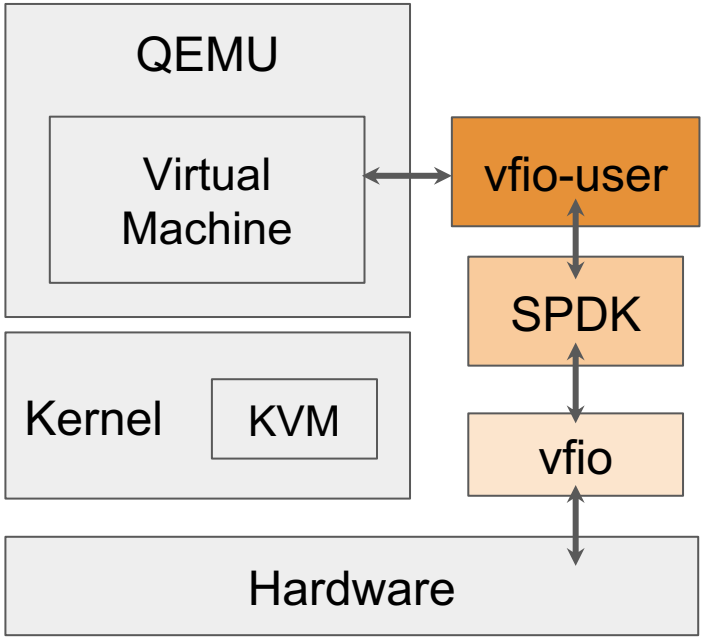

McGill UNIVERSITY

I want my process to see an NVMe

McGill UNIVERSITY

Fast and simple

McGill UNIVERSITY

Fast, involves some configuration

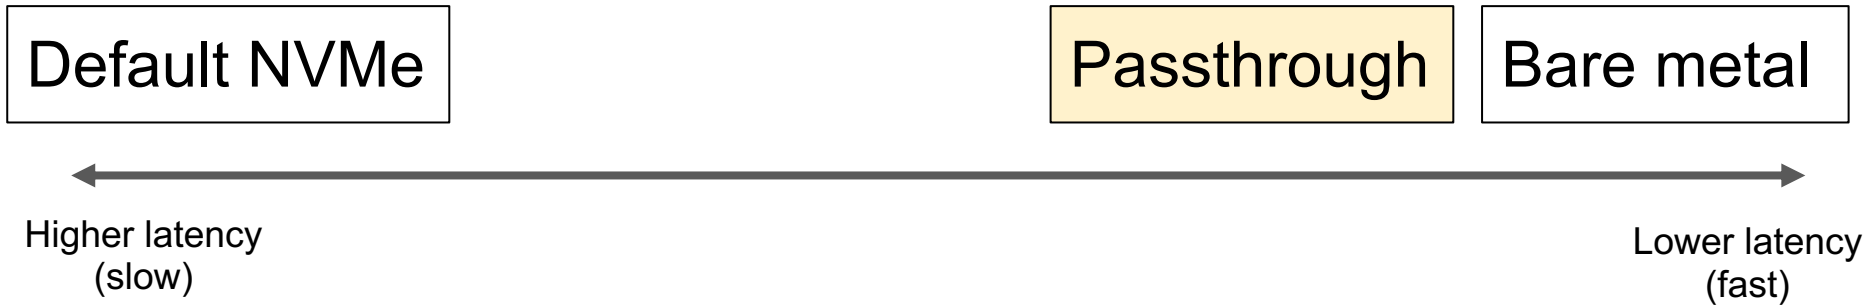

McGill UNIVERSITY

Slow, requires one command-line flag

McGill UNIVERSITY

Where does vfio-user fit?

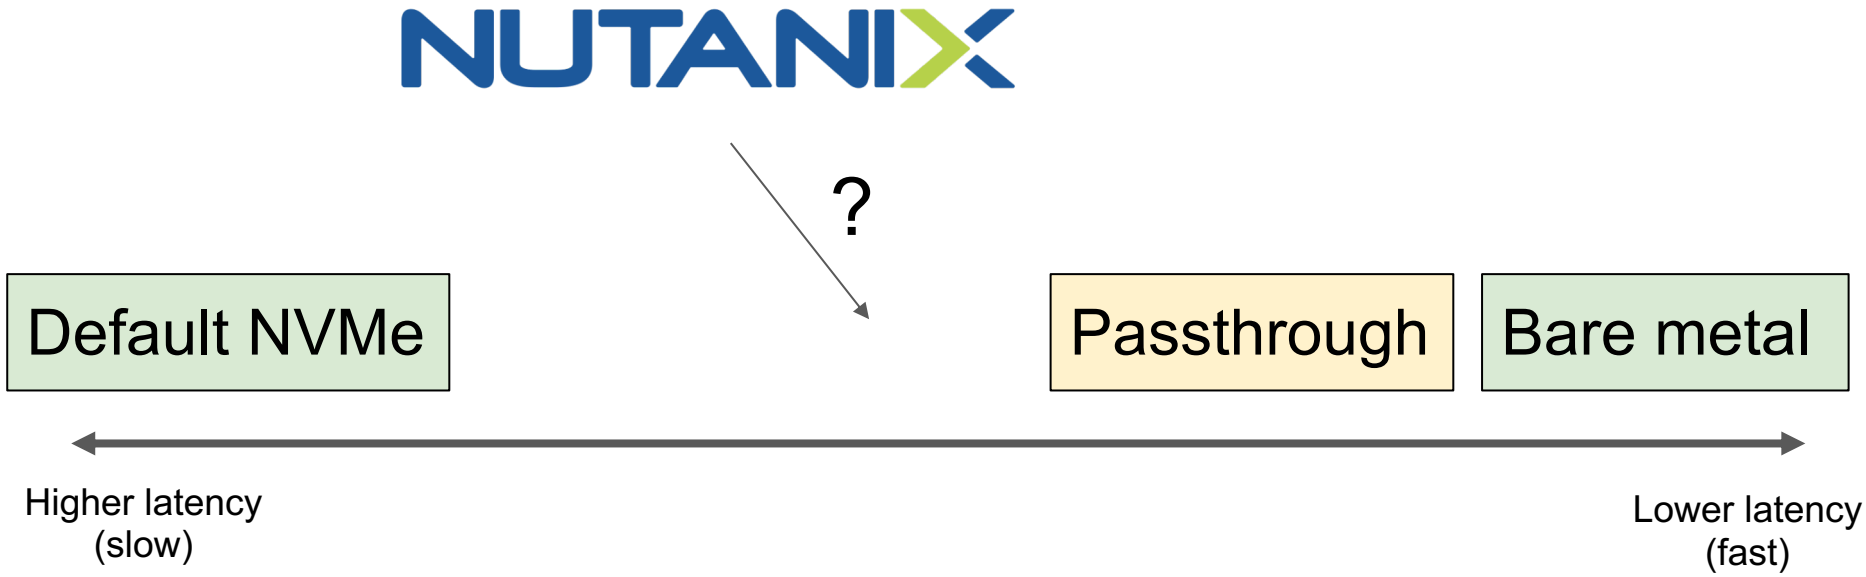

McGill UNIVERSITY

Bare metal configuration

Default | Passthrough | Bare metal

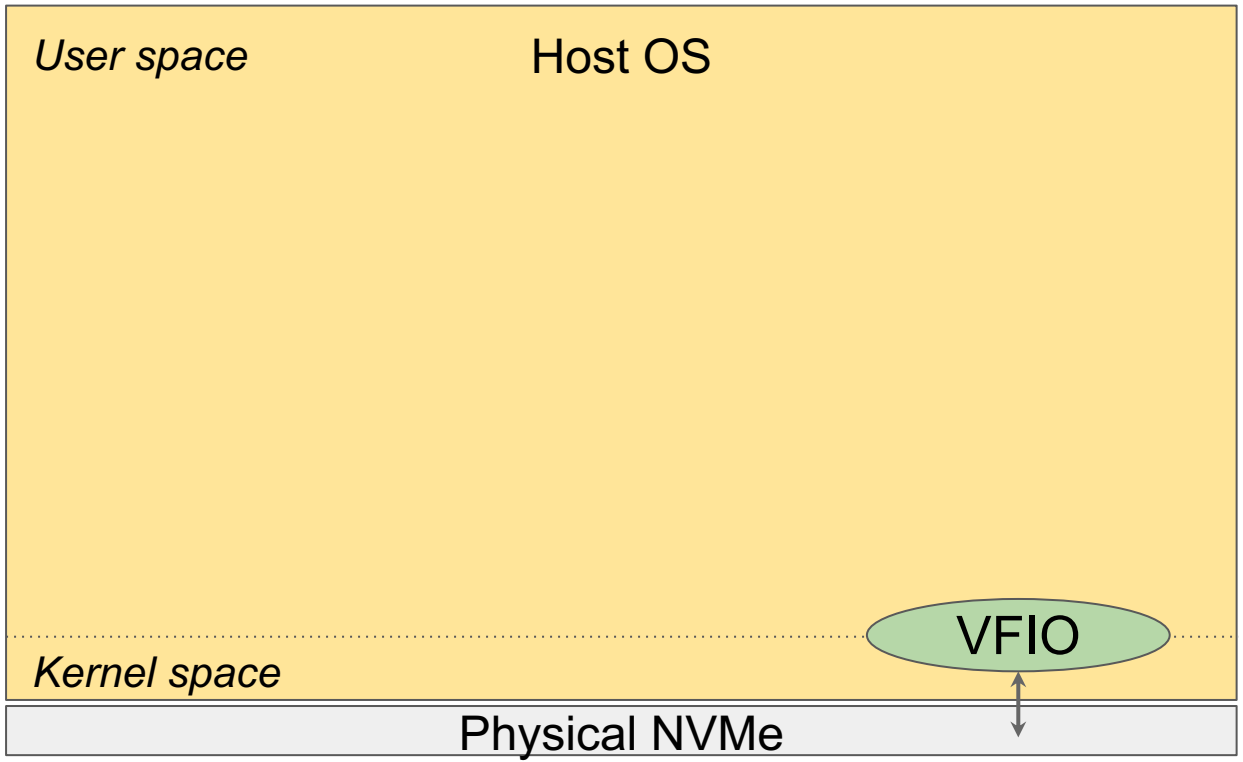

McGill UNIVERSITY

Bare metal configuration

Default — Passthrough — Bare metal

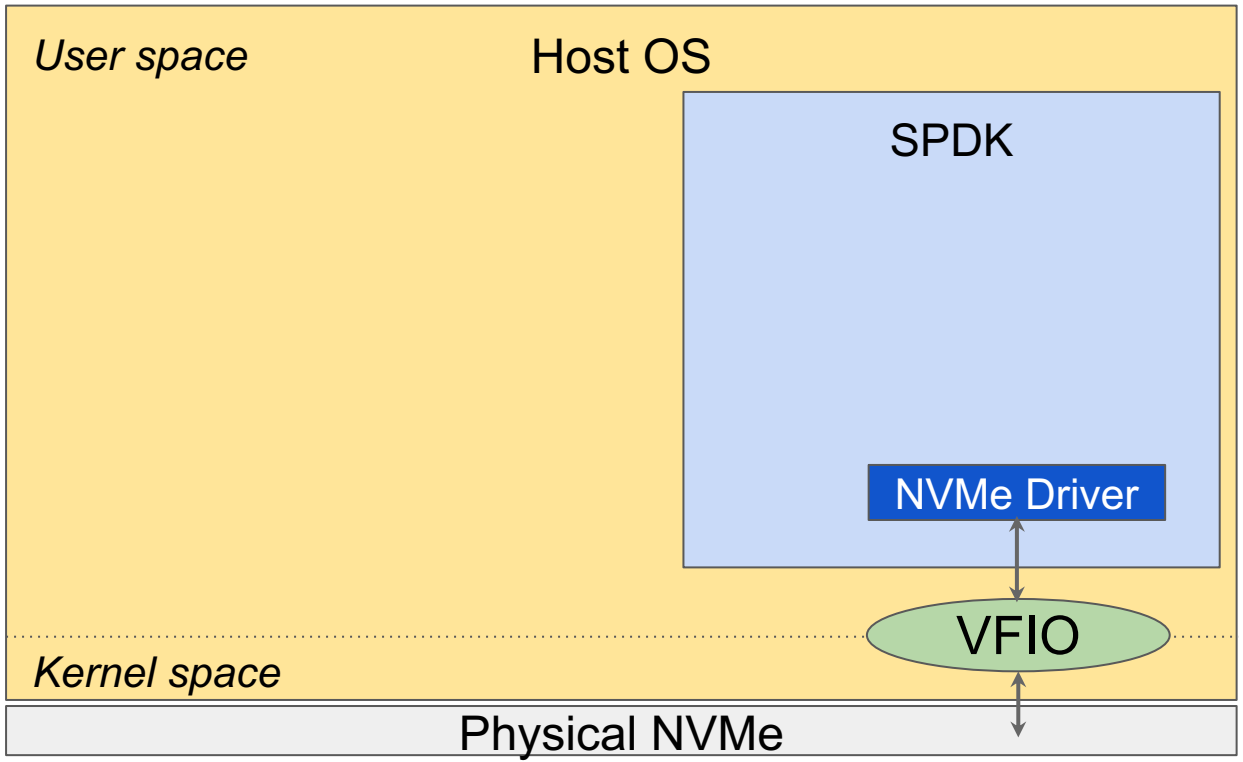

McGill UNIVERSITY

Bare metal configuration

Default | Passthrough | Bare metal

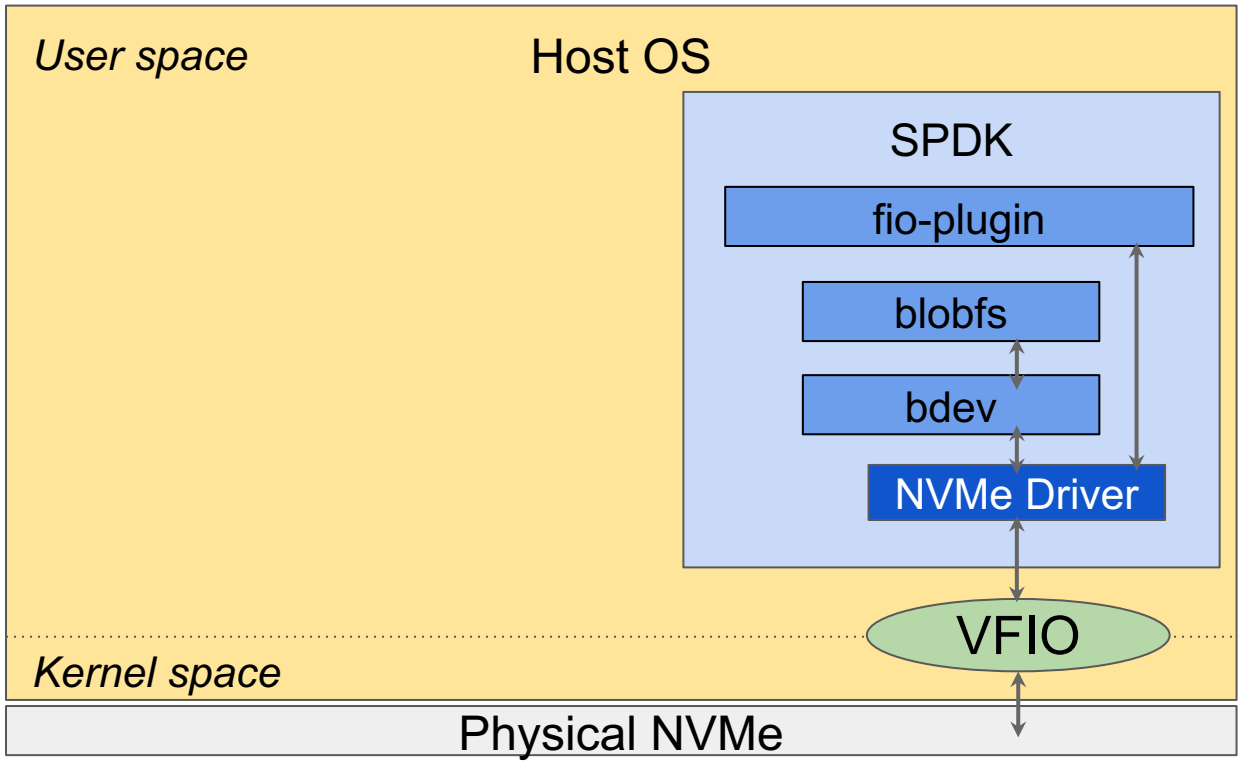

McGill UNIVERSITY

Bare metal configuration

Default | Passthrough | Bare metal

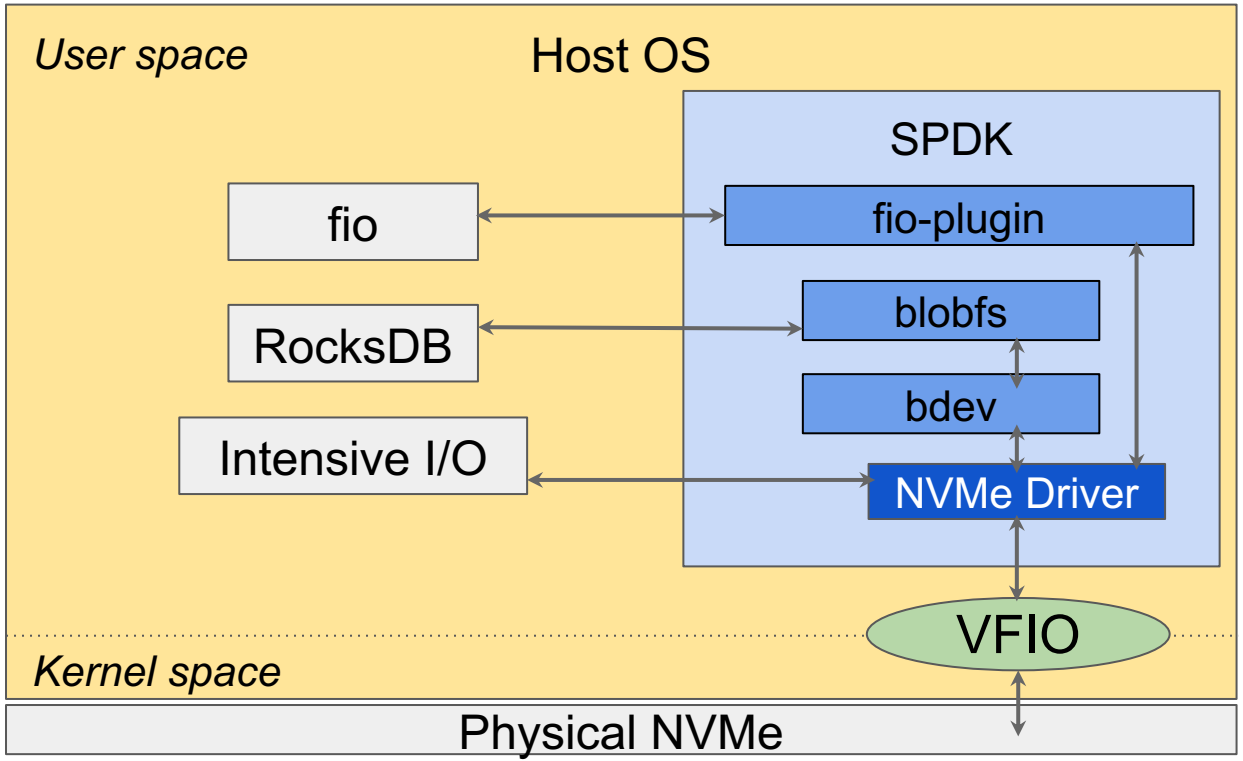

McGill UNIVERSITY

Default NVMe configuration

Default | Passthrough | Bare metal

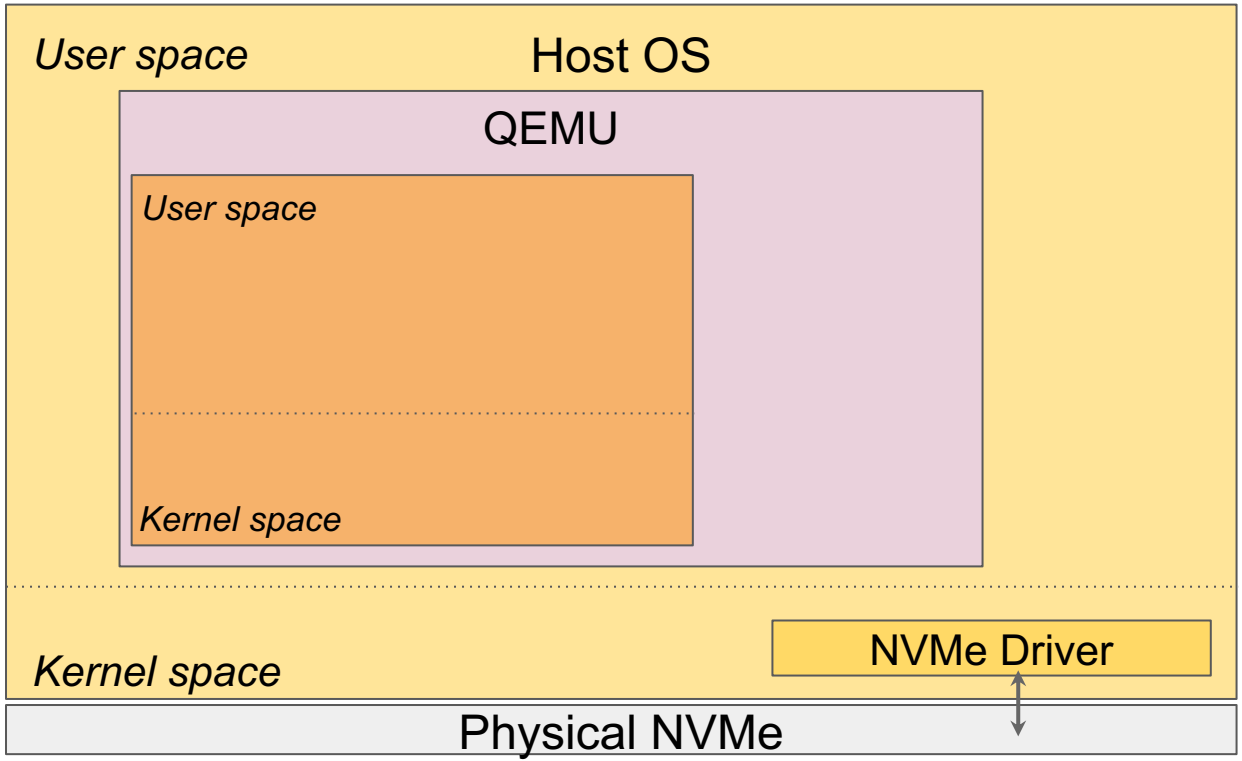

McGill UNIVERSITY

Default NVMe configuration

Default | Passthrough | Bare metal

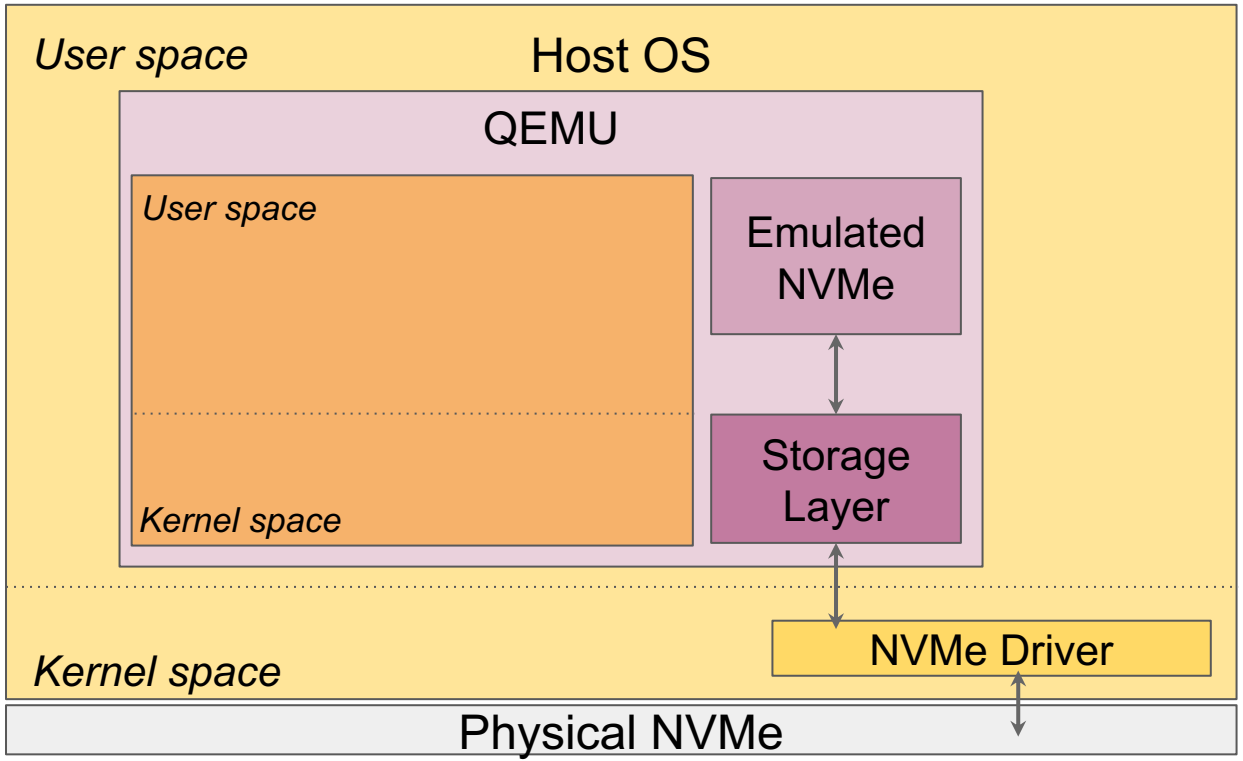

McGill UNIVERSITY

Default NVMe configuration

Default | Passthrough | Bare metal

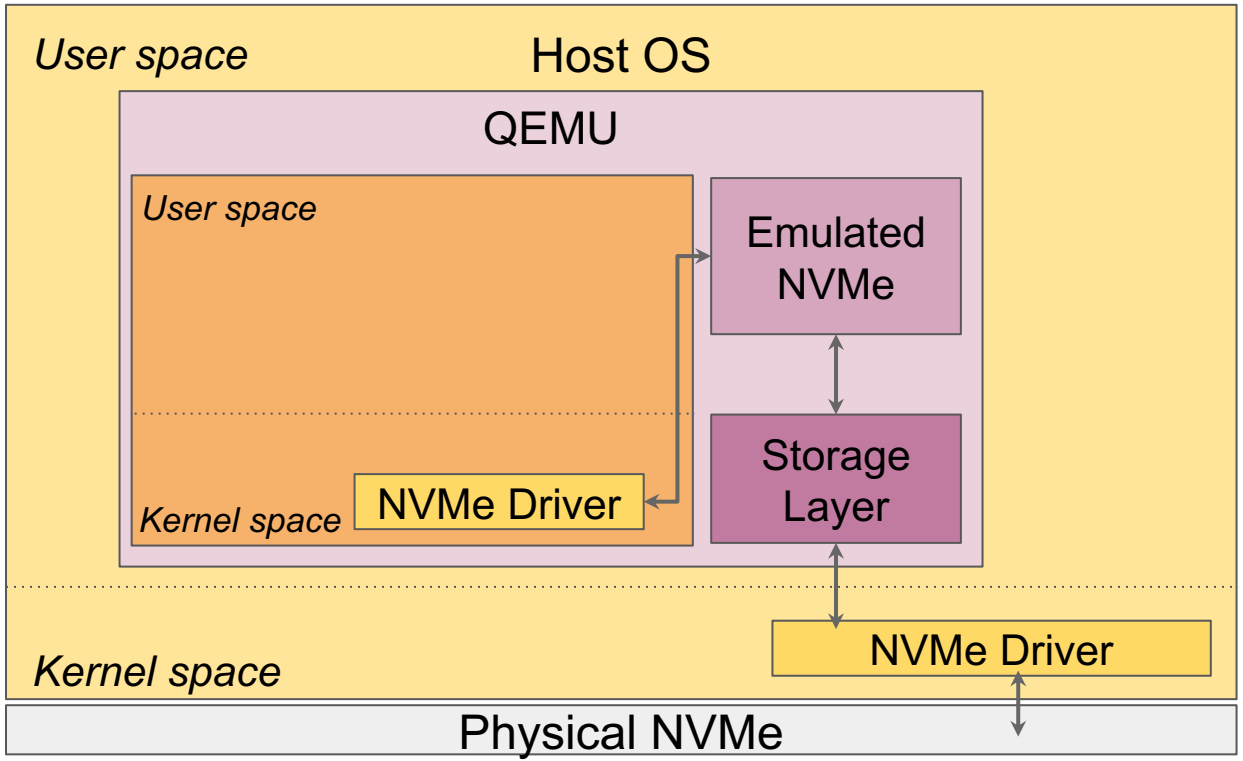

McGill UNIVERSITY

Default NVMe configuration

Default | Passthrough | Bare metal

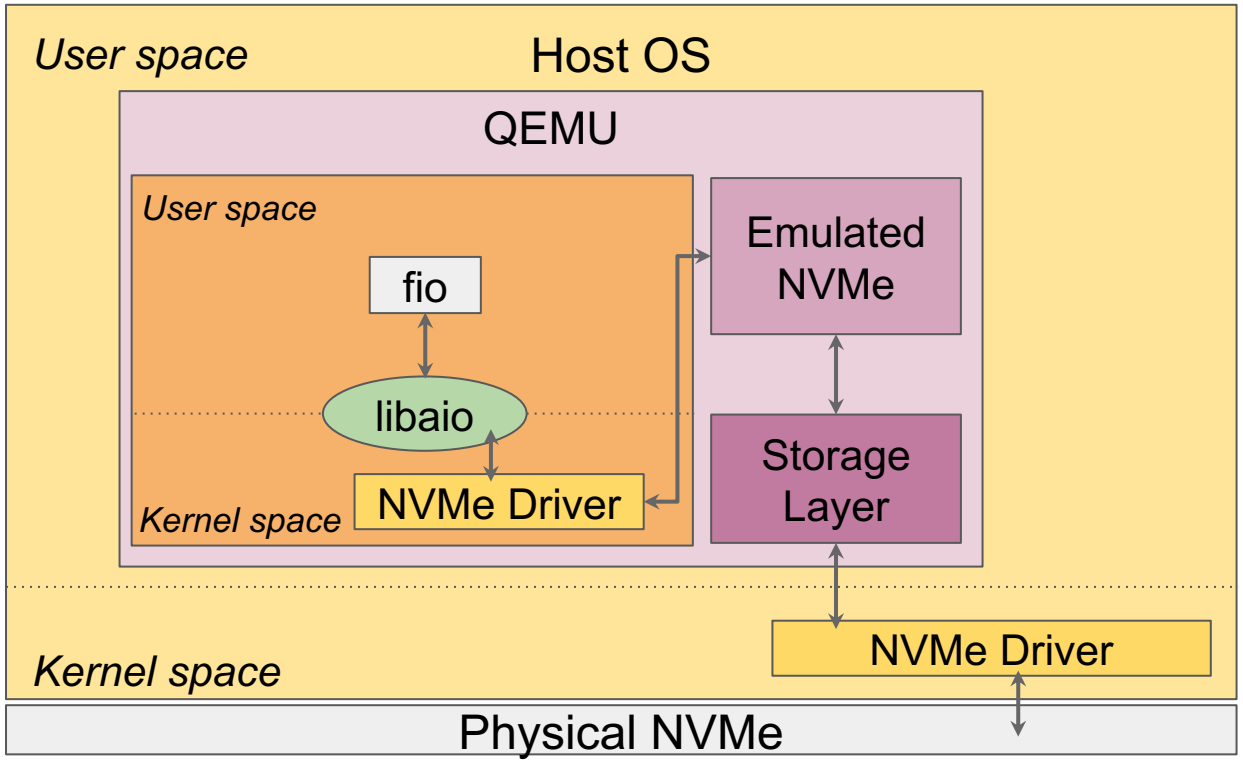

McGill UNIVERSITY

Passthrough configuration

Default | Passthrough | Bare metal

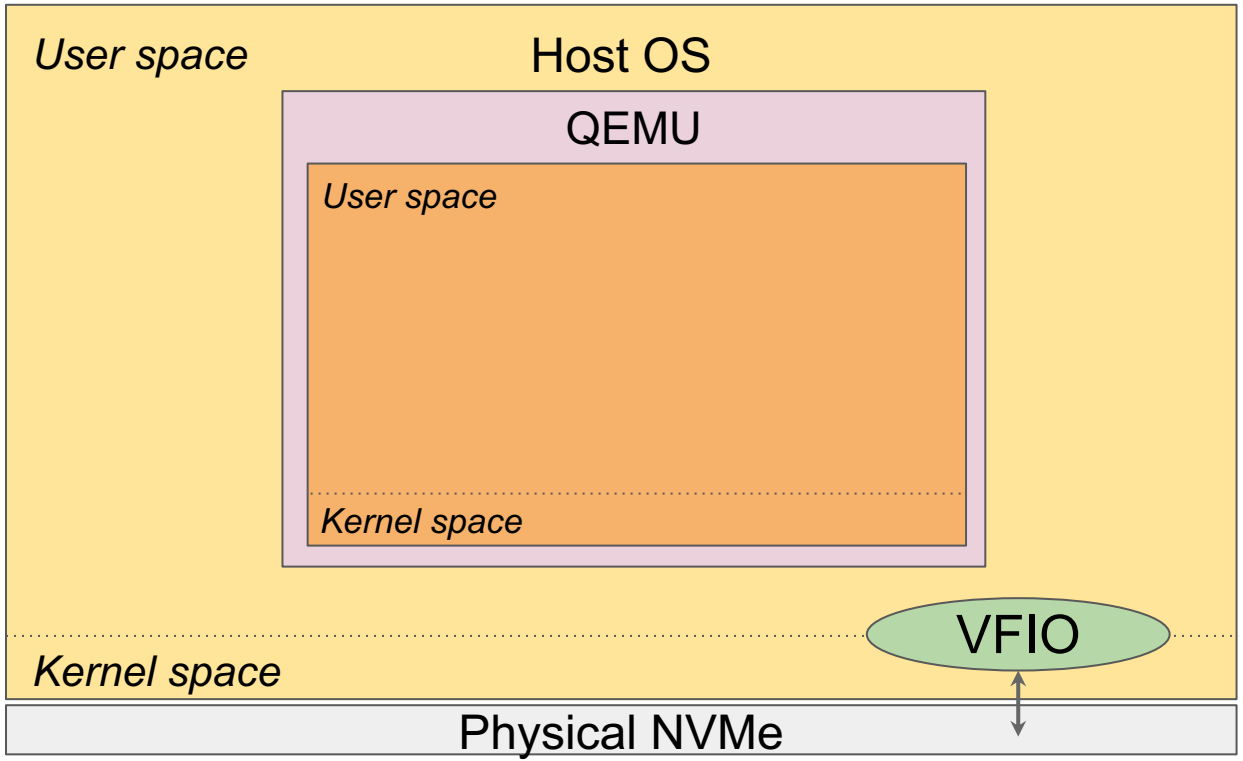

McGill UNIVERSITY

Passthrough configuration

Default | Passthrough | Bare metal

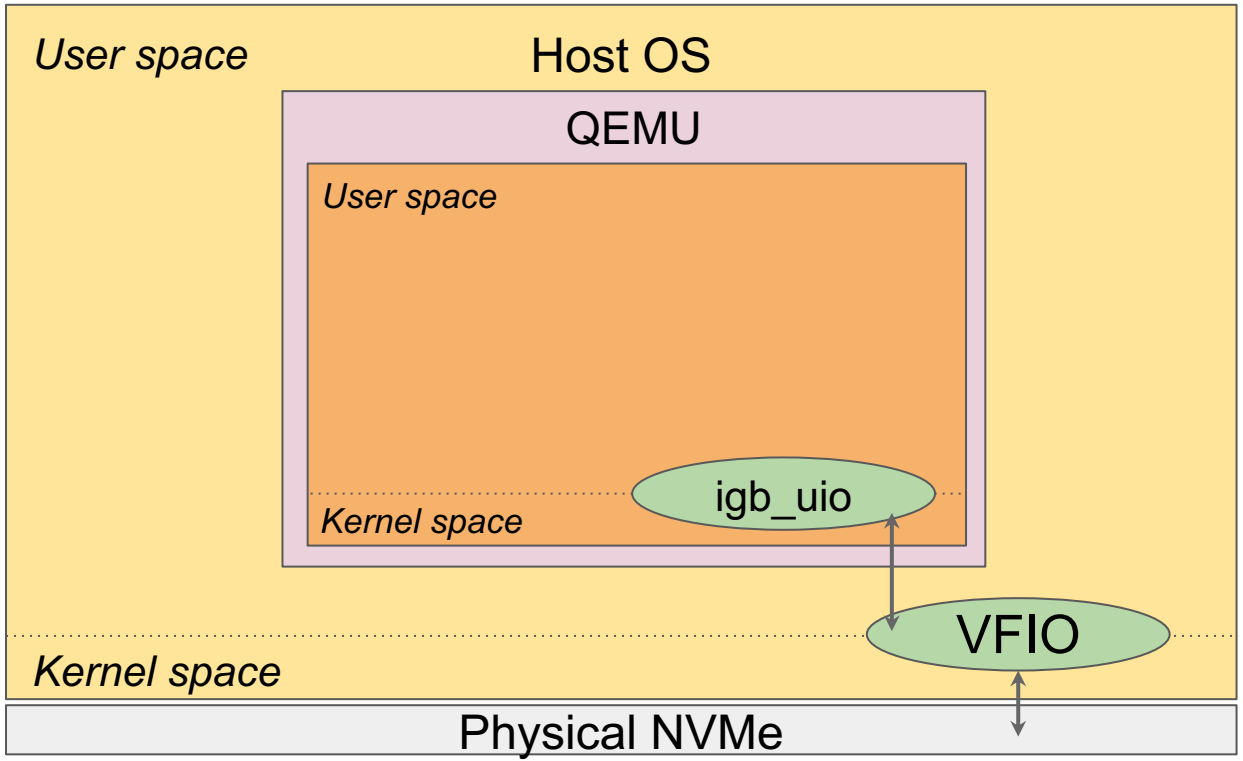

McGill UNIVERSITY

Passthrough configuration

Default ⟷ Passthrough | Bare metal

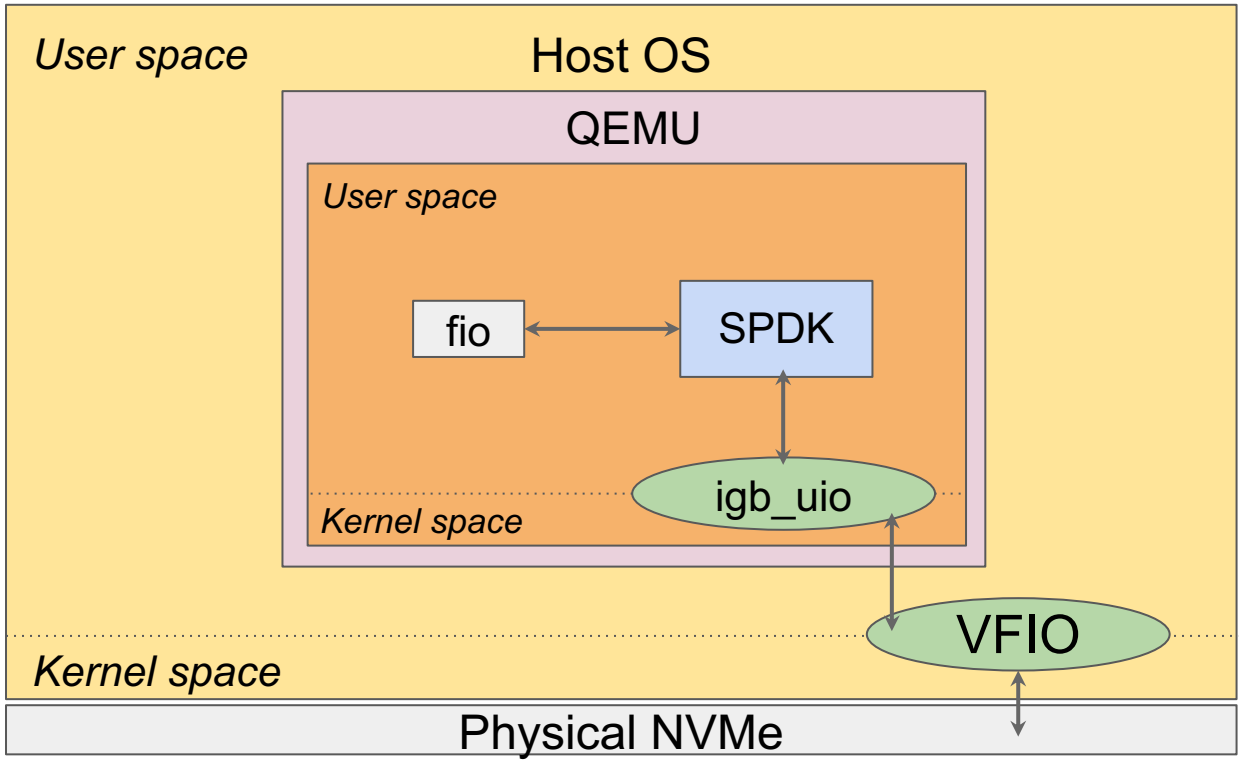

McGill UNIVERSITY

Vfio-user configuration

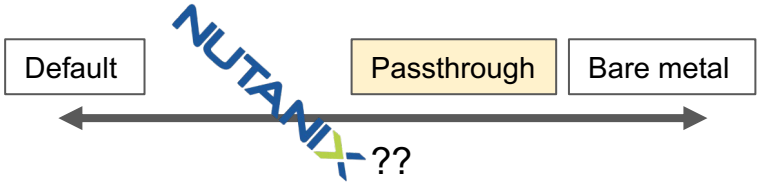

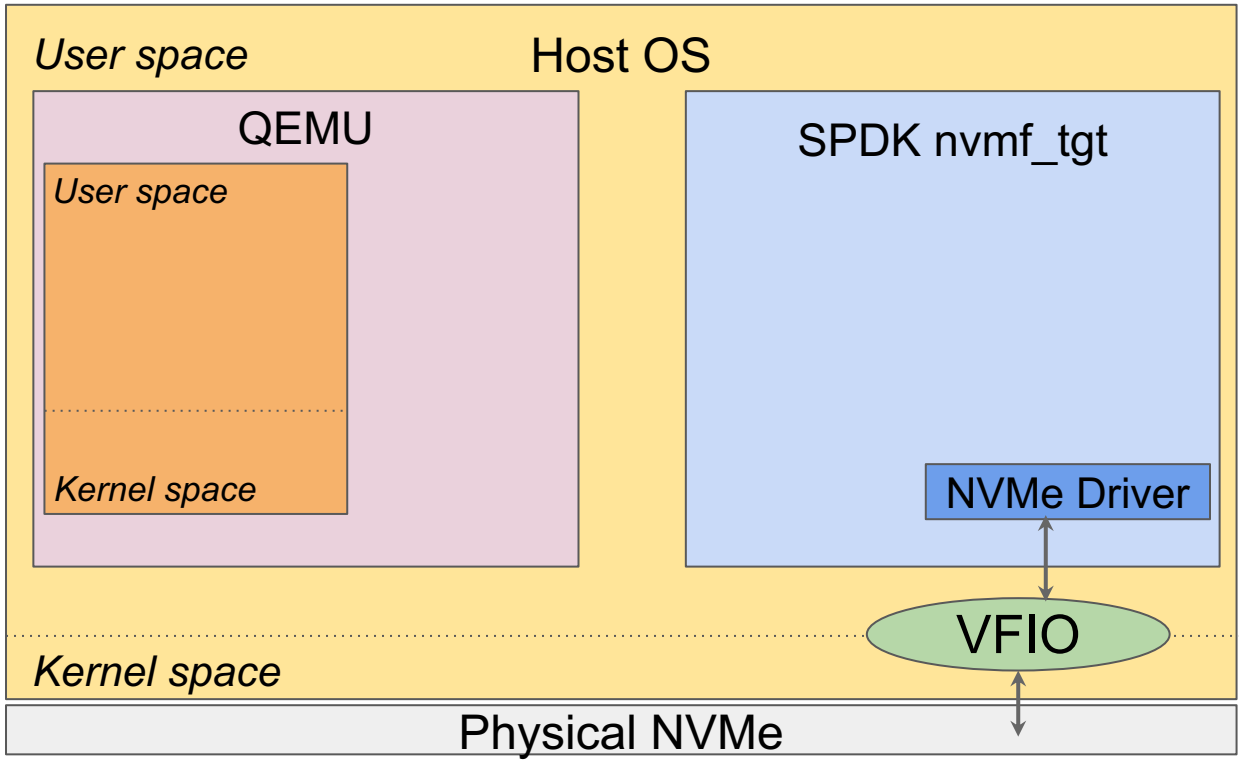

McGill UNIVERSITY

Vfio-user configuration

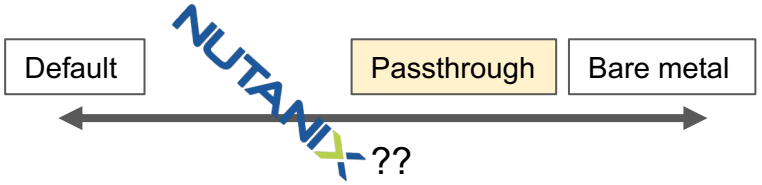

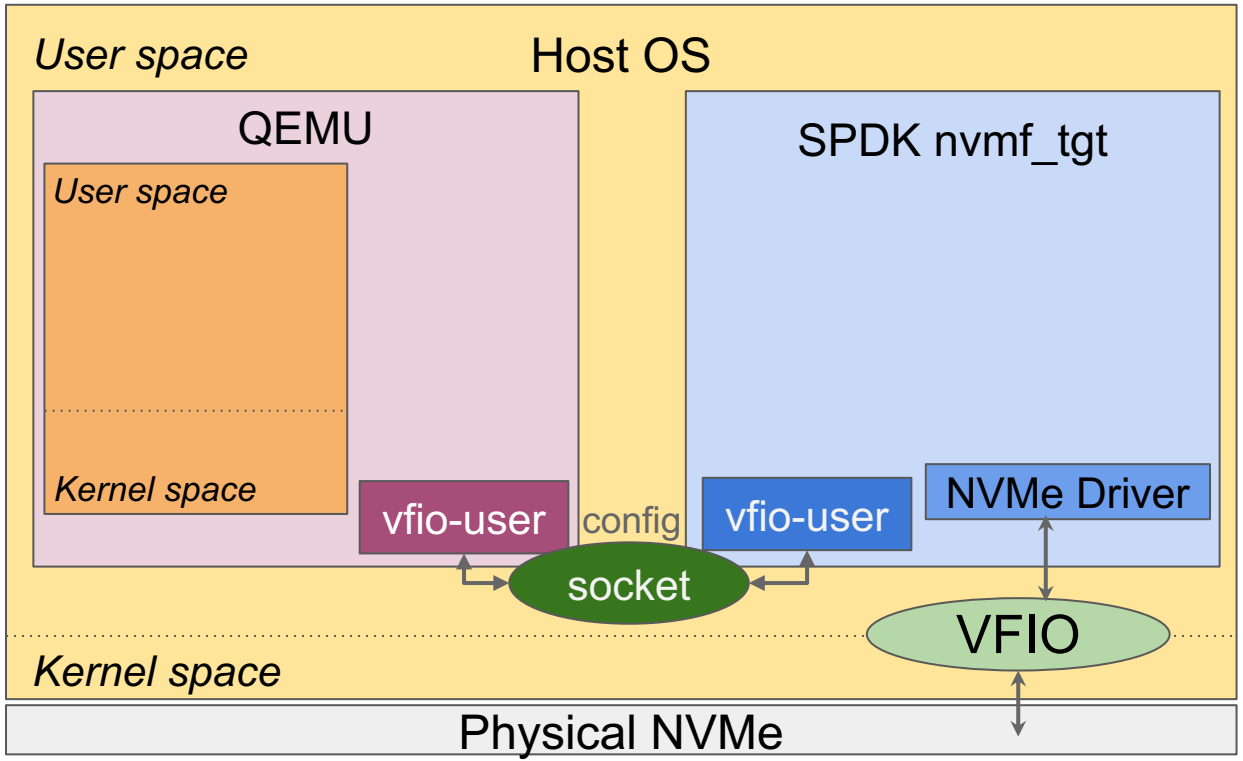

Vfio-user configuration

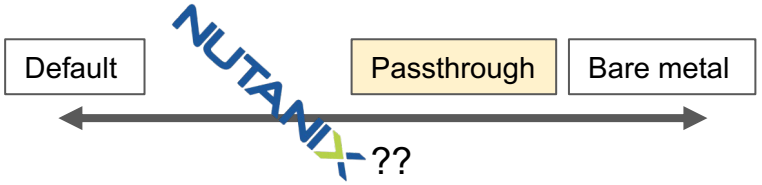

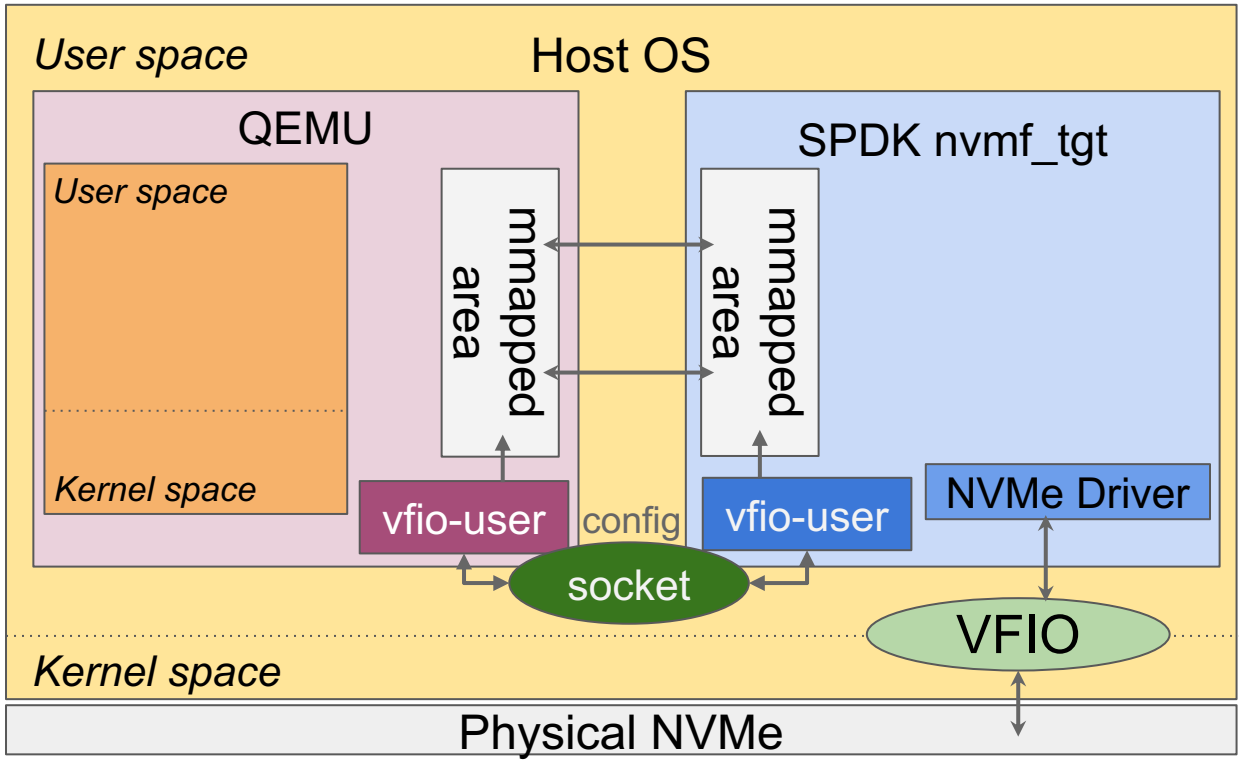

McGill UNIVERSITY

Vfio-user configuration

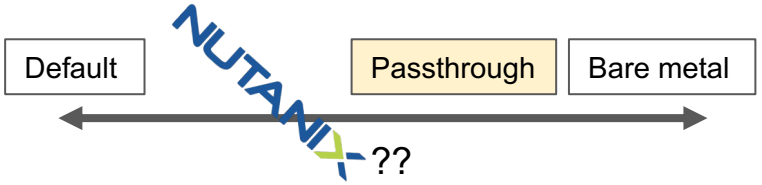

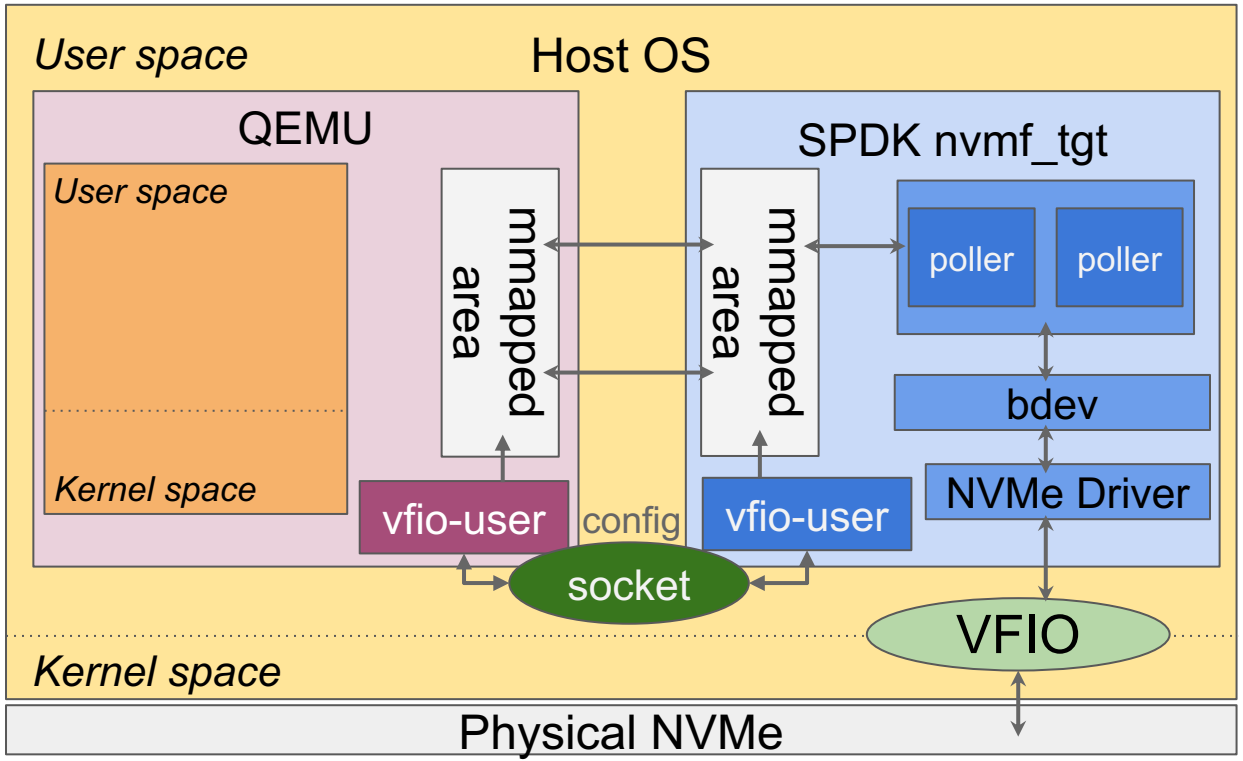

Vfio-user configuration

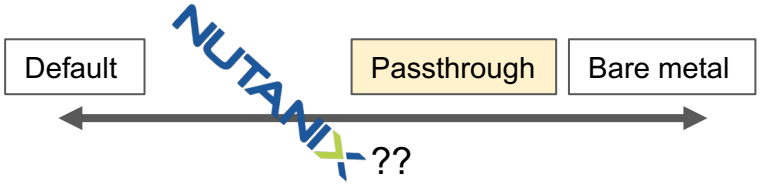

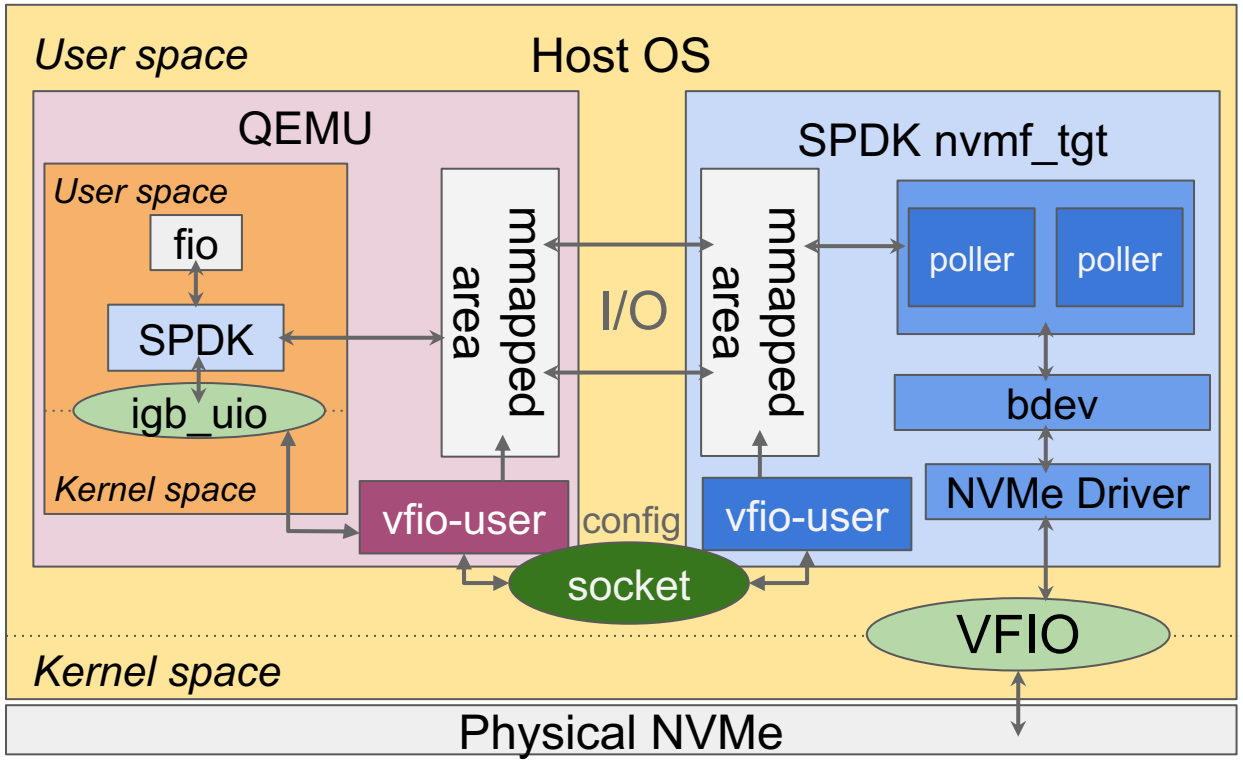

Experimental evaluation

Hardware Specifications	
CPU	36 Core Xeon Gold 6240L @ 2.40 GHz
Memory	768 GiB 3200MHz DDR4 DIMM
SSD	375GiB Dell Express Flash NVMe P4800X

McGill UNIVERSITY

Results: fio random reads

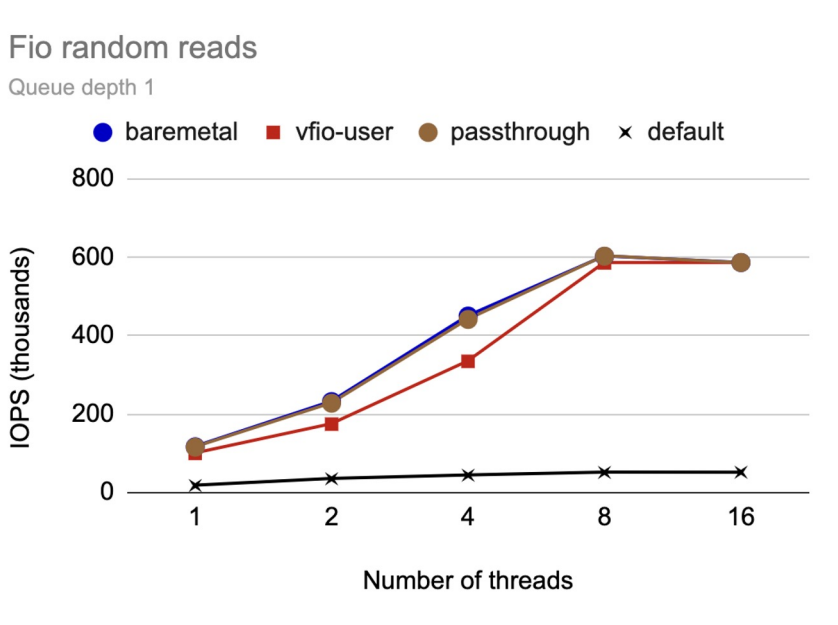

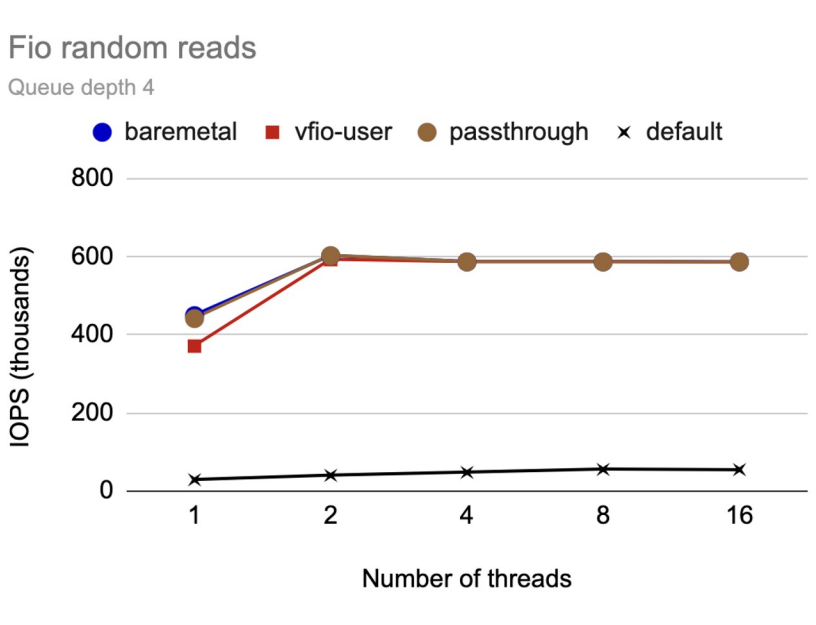

McGill UNIVERSITY

Results: fio random writes

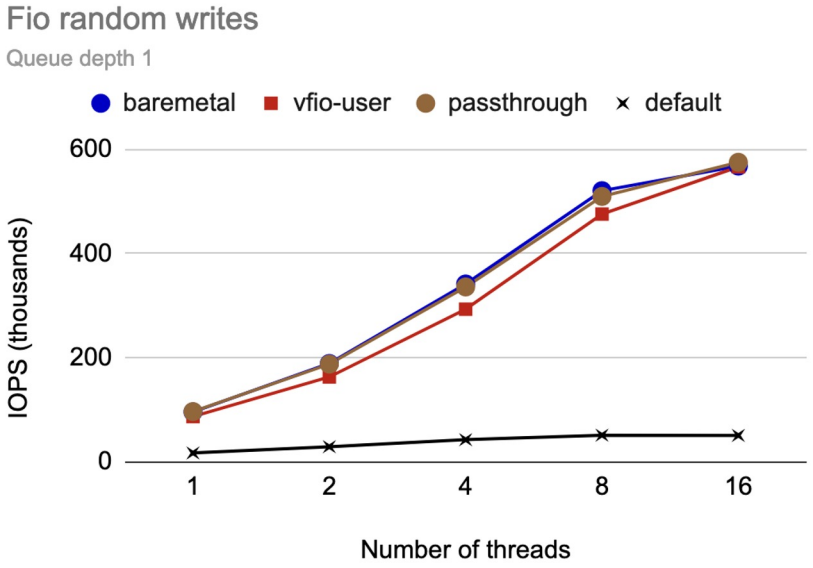

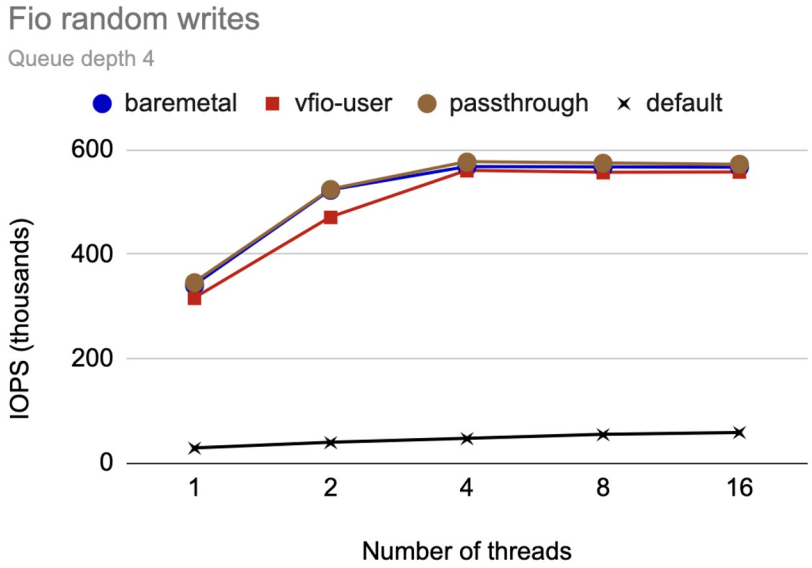

McGill UNIVERSITY

Results: RocksDB benchmarks

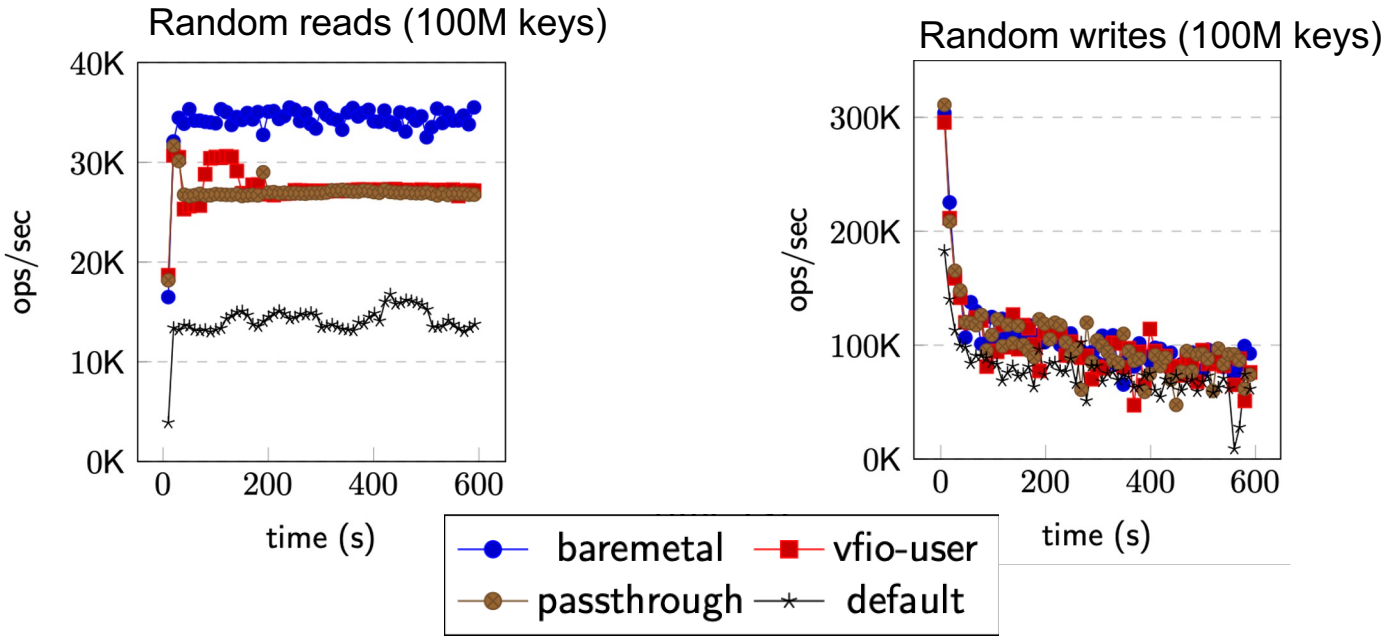

McGill UNIVERSITY

Layer-by-layer latency measurement

Replaced internal SPDK with synchronous I/O and guest NVMe driver

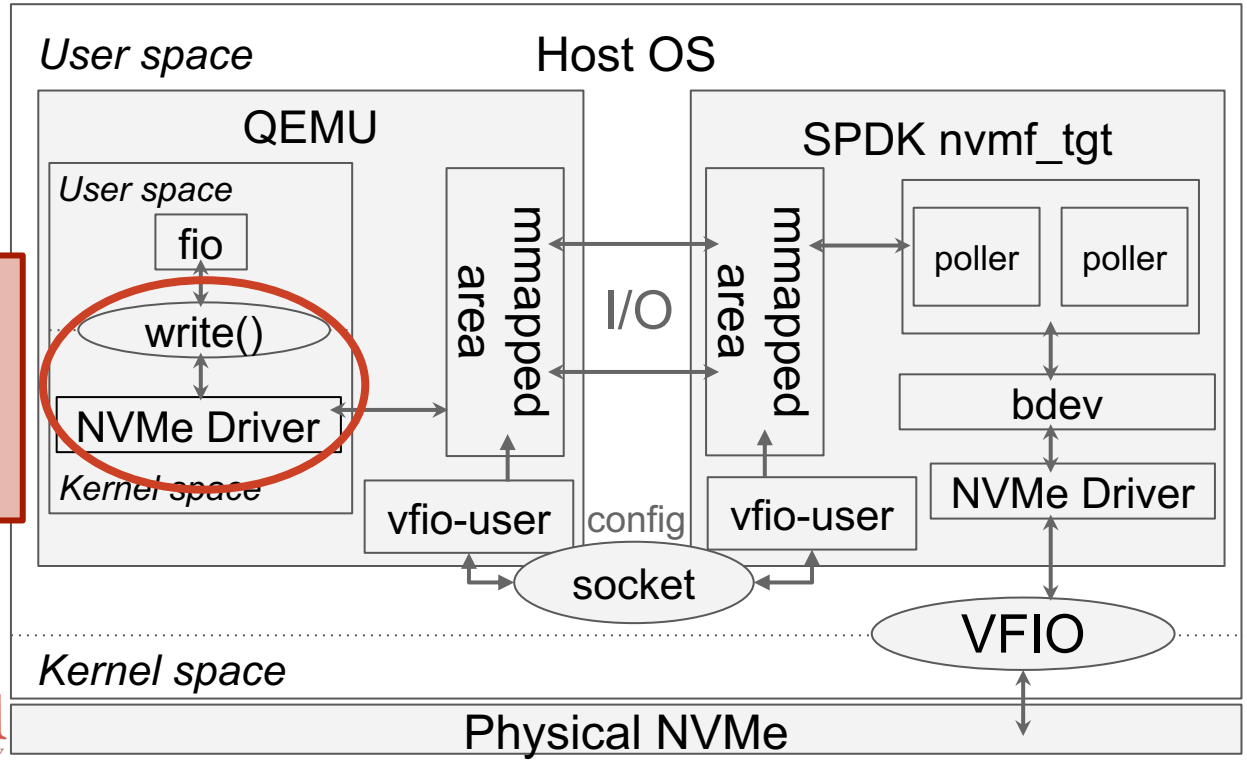

Layer-by-layer latency measurement

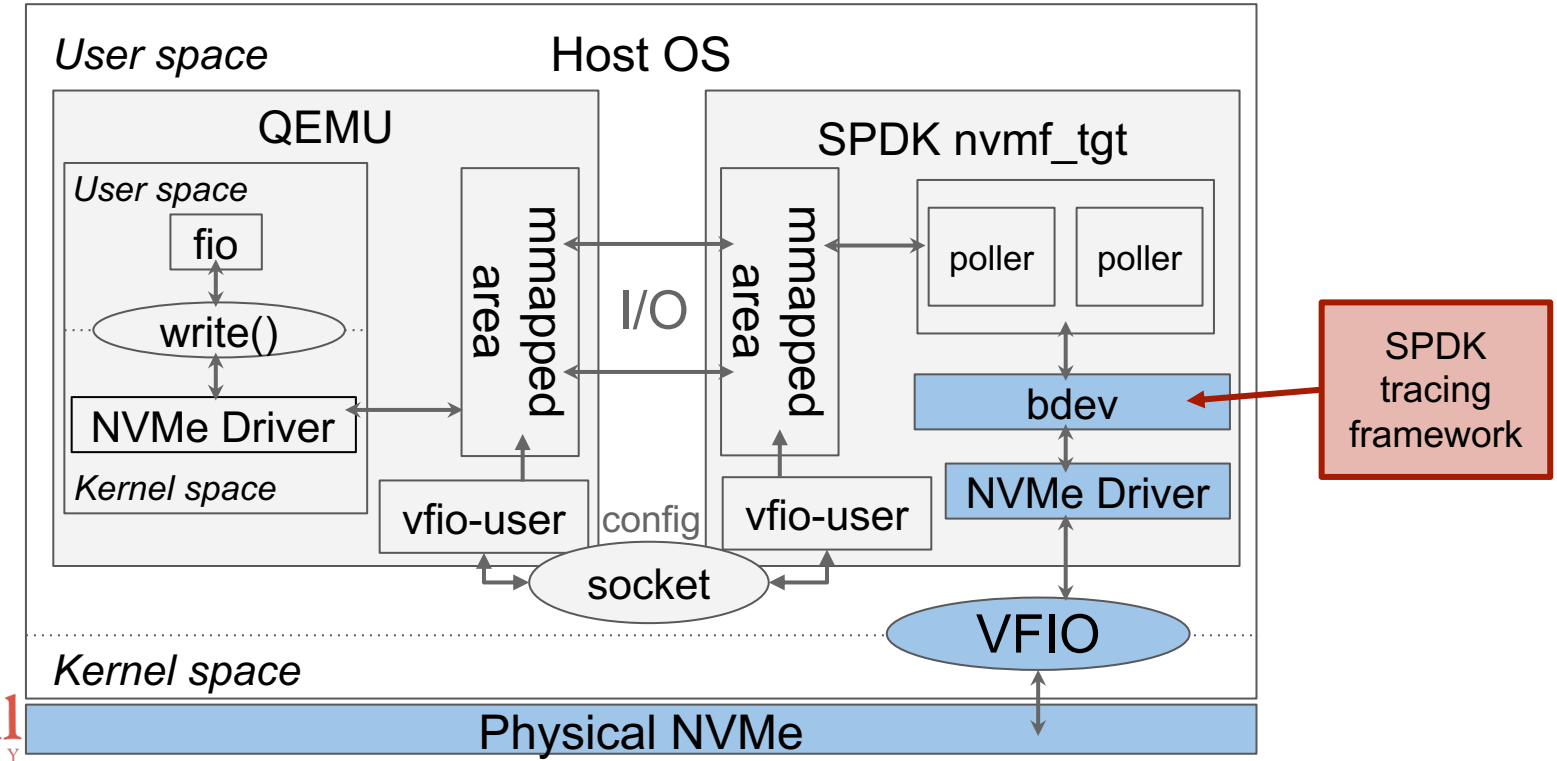

Layer-by-layer latency measurement

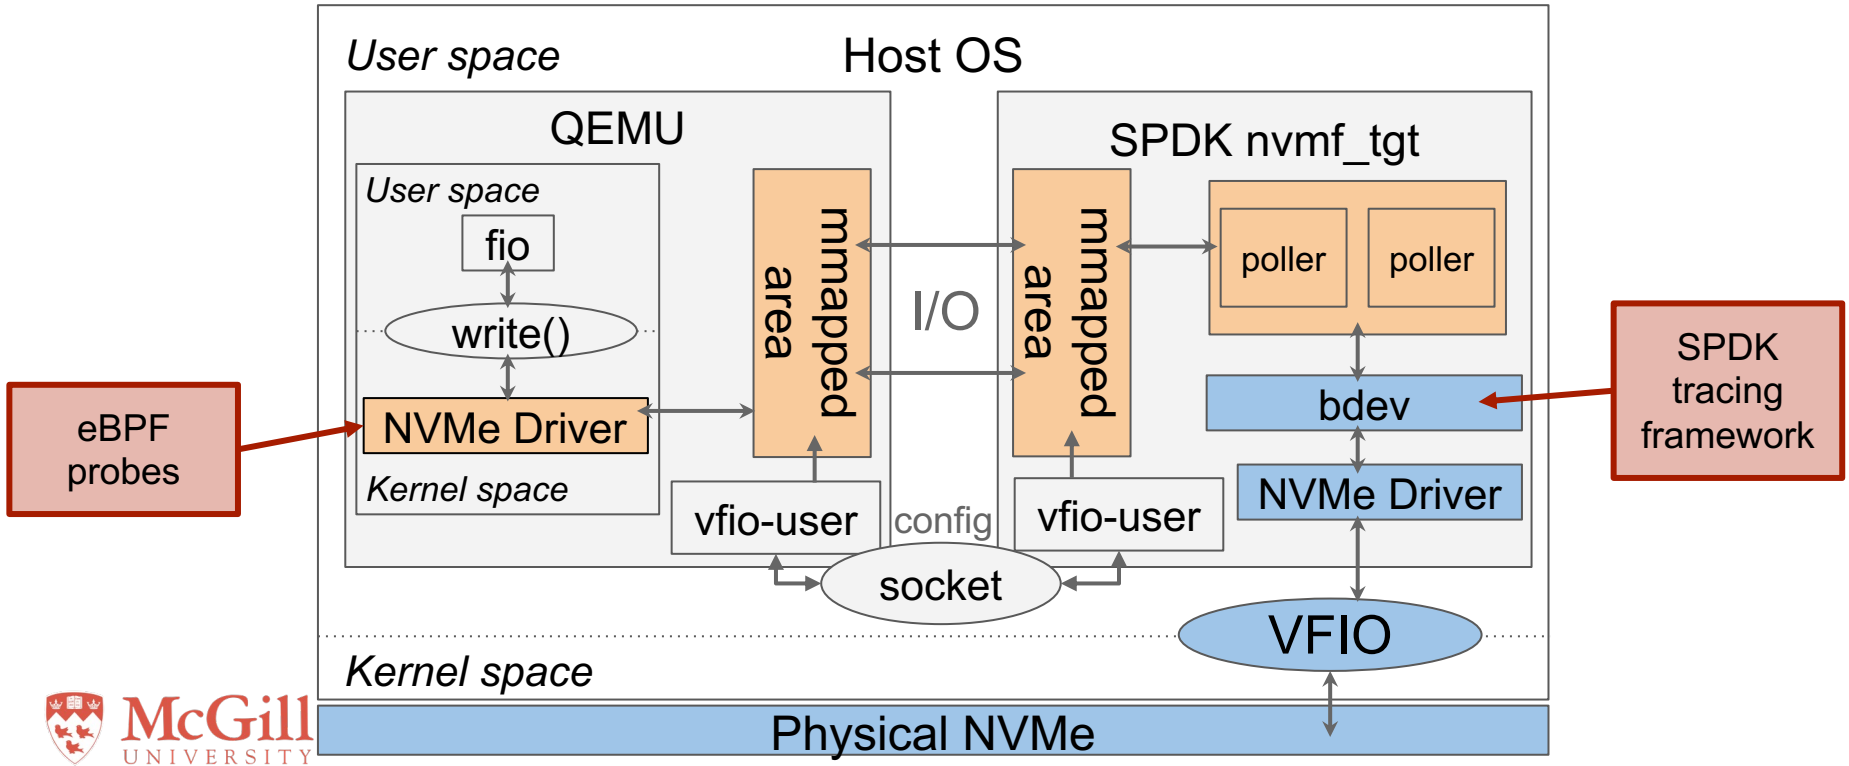

Layer-by-layer latency measurement

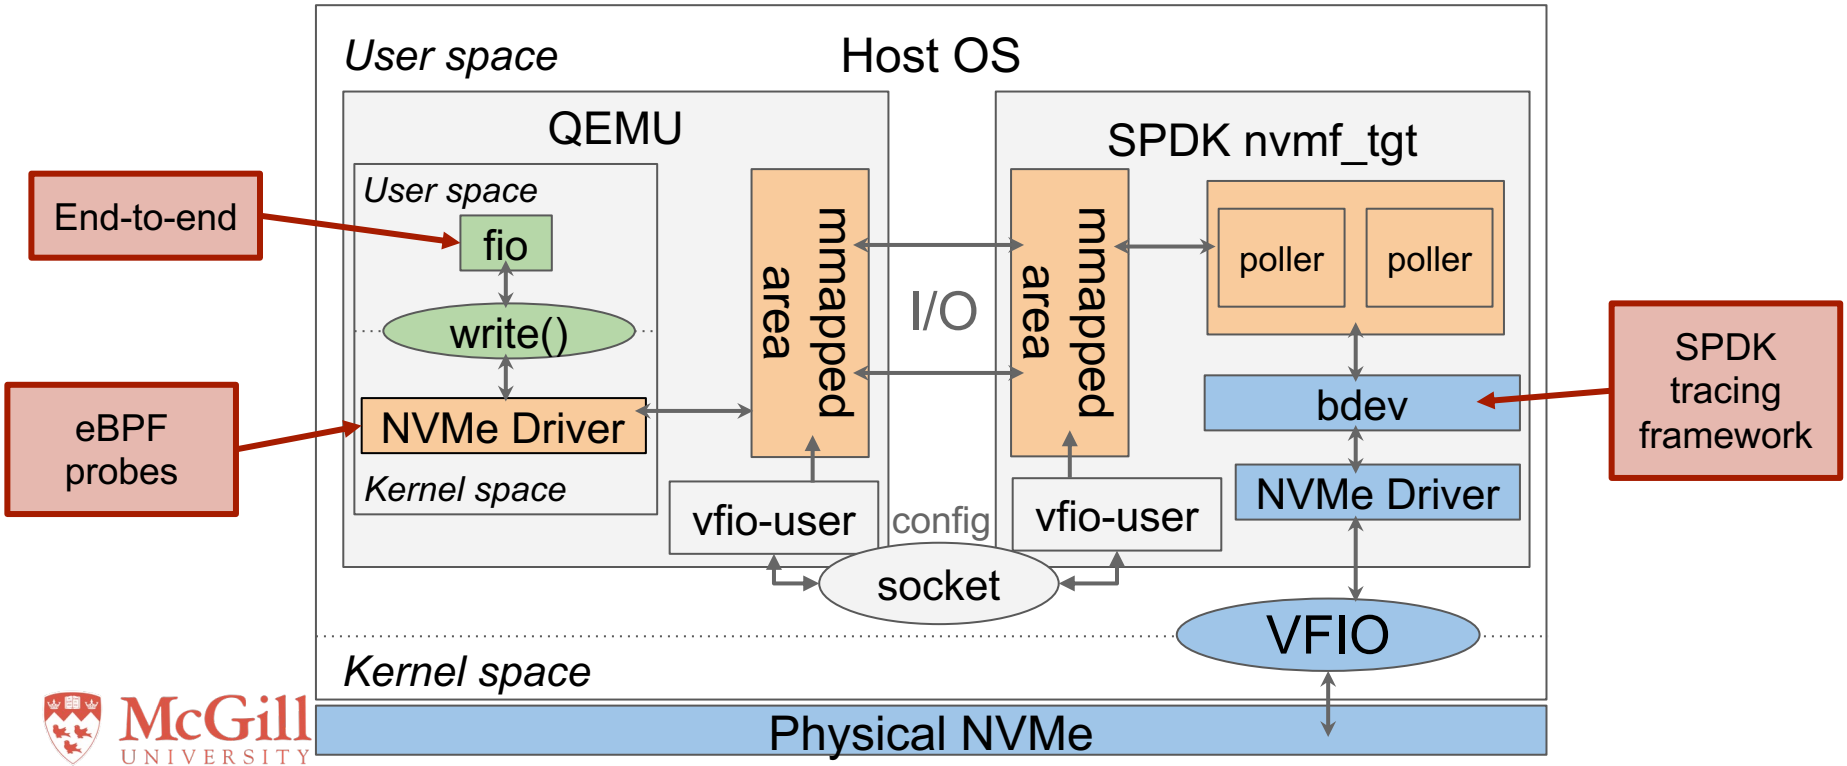

Layer-by-layer latency measurements

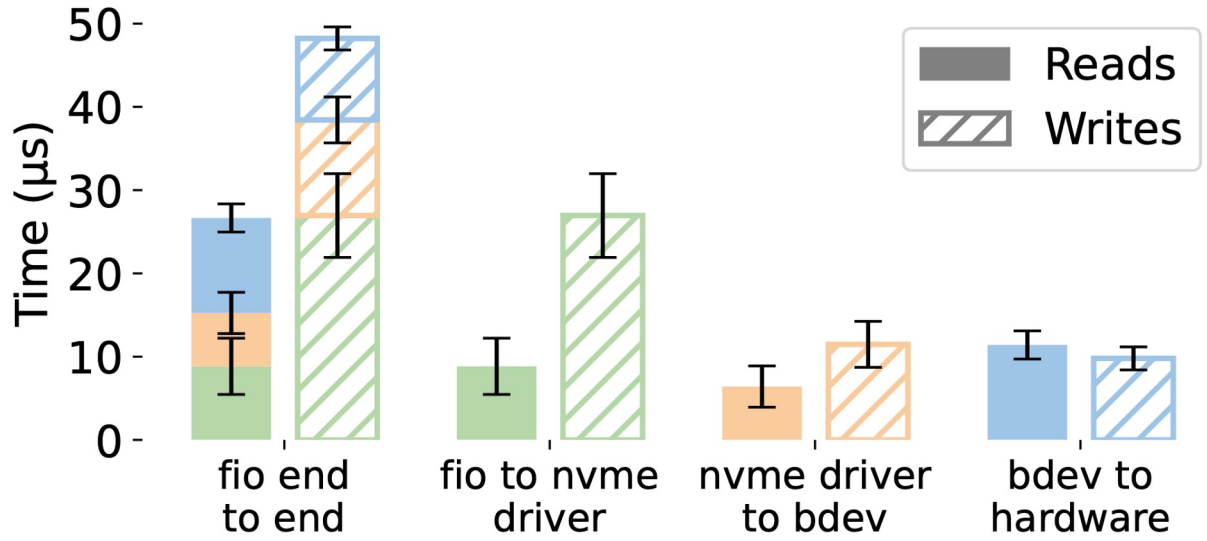

McGill UNIVERSITY

Where does vfio-user fit again?

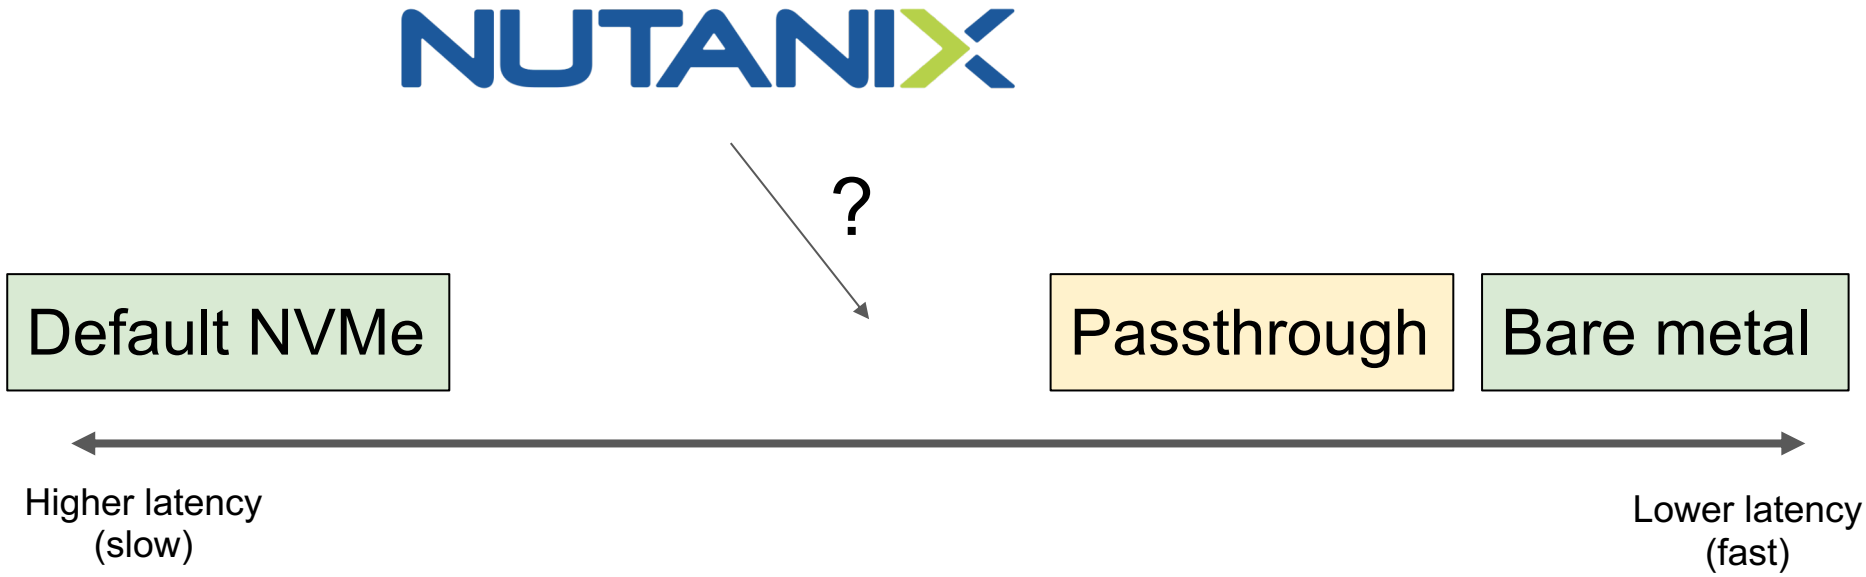

McGill UNIVERSITY

Right about next to Passthrough

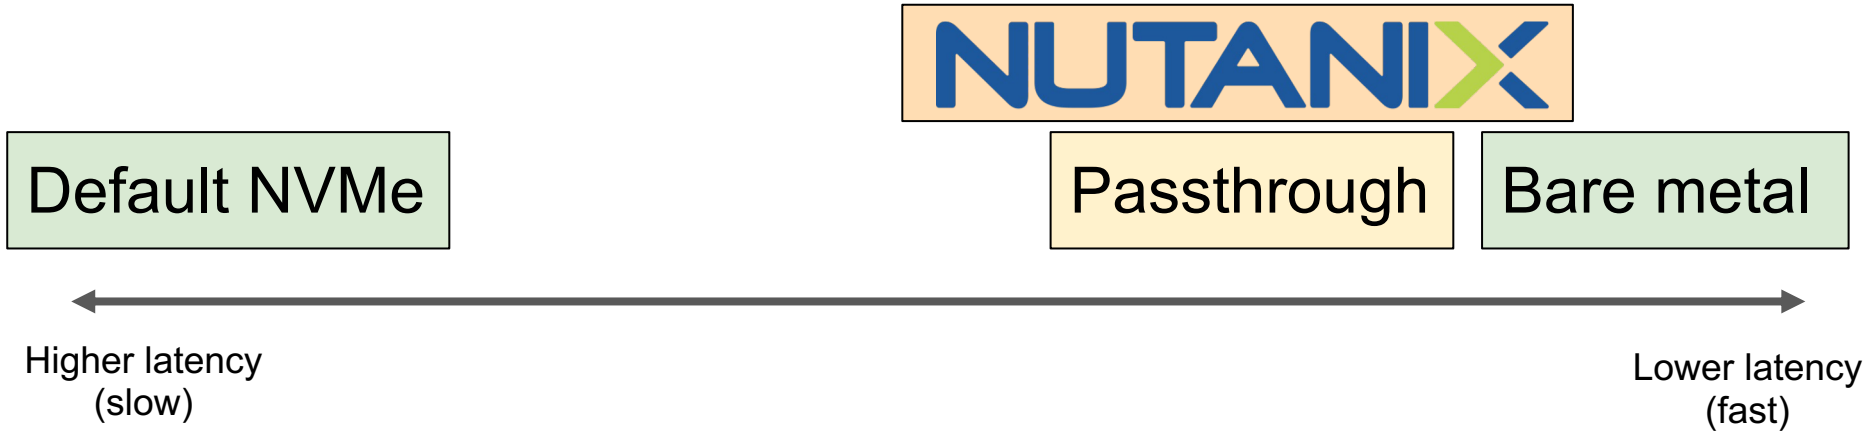

McGill UNIVERSITY

Conclusion

- **Vfio-user** appears to have comparable performance to passthrough
- Could be viable for VM storage
- See more benchmarks and analysis in our paper!

Visit the DISCS Lab:
https://discslab.cs.mcgill.ca

On the job market!
js@rolon.co

McGill UNIVERSITY