Mellanox Bifurcated DPDK PMD

Rony Efraim

DPDK summit Dublin Oct 2016

Mellanox TECHNOLOGIES

Connect. Accelerate. Outperform.™

Mellanox PMD

- Mellanox PMD is built on top of **libibverbs** using the Raw Ethernet Accelerated Verbs API.
- Verbs is standard API of openfabrics.org
- The PMD directly access the NIC for fast data path
- Control path is going through the kernel modules.
- Unlike other PMDs, Mellanox PMD is yet another user space application and the NIC is still being controlled by the kernel

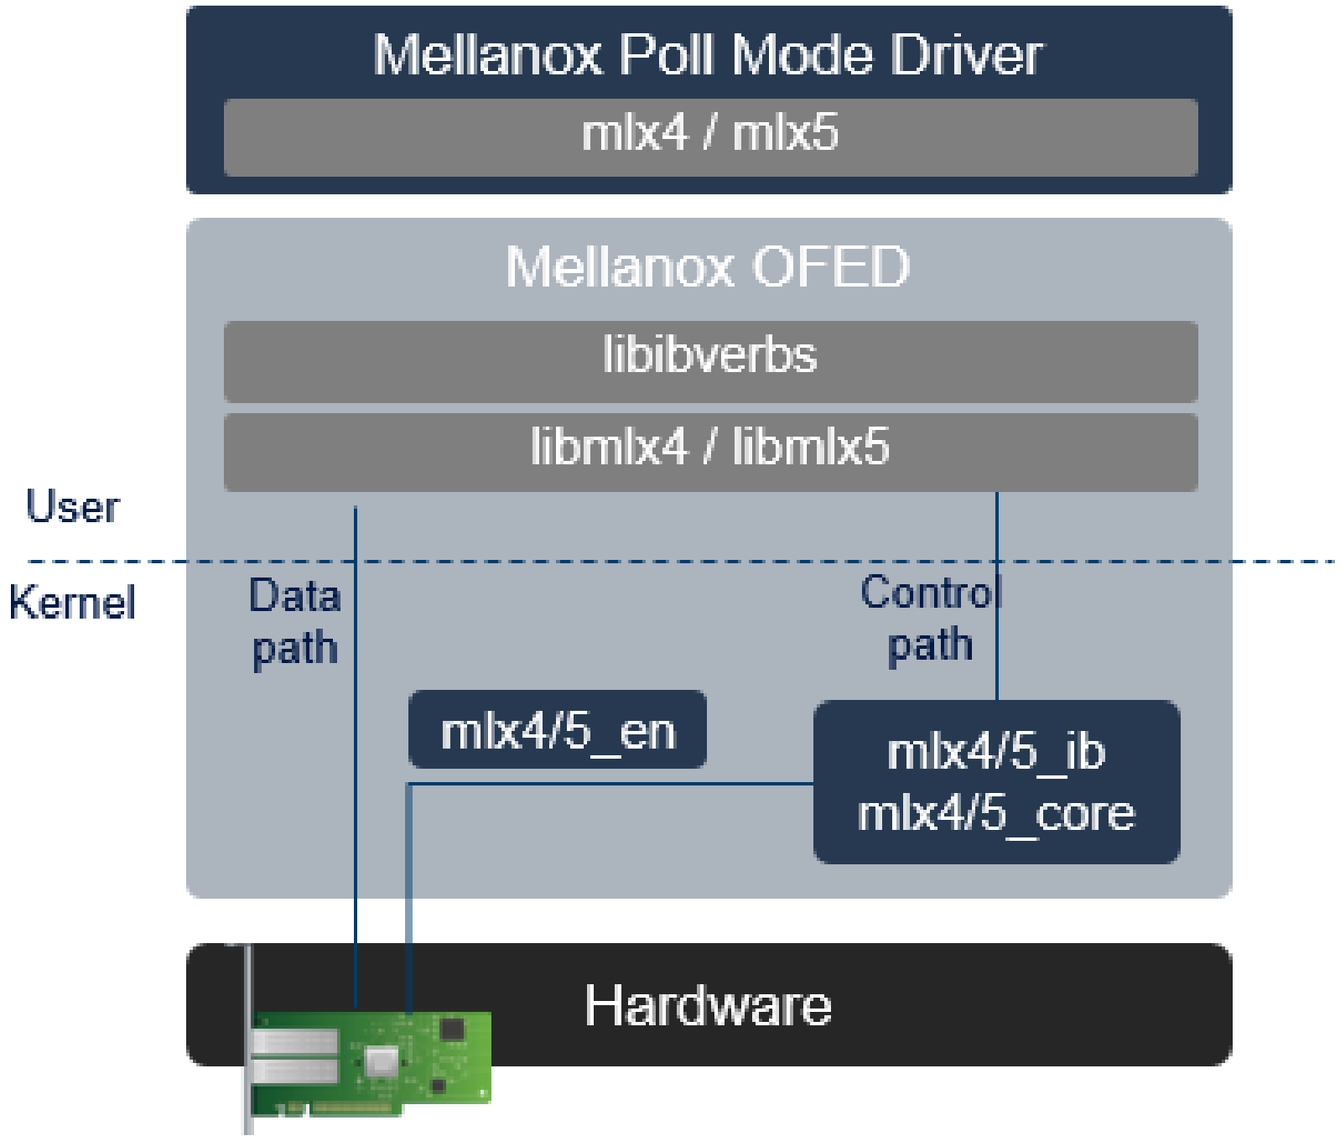

Kernel driver

- Control commands (e.g ethtool) will work as usual
- Mellanox PMD relies on system calls for control operations such as querying/updating the MTU and flow control parameters.
- For security reasons and robustness, the PMD only deals with virtual memory addresses.
- DPDK steer the required traffic by using DPDK filter.
- All the leftover is going to the kernel driver.
- The kernel handle the leftover traffic and can replay for arp, ping ….

Multiple DPDK Applications on single device

- You can run multiple application on a single function.
- resources allocations are handled by the kernel combined with hardware specifications that allow it to handle virtual memory addresses directly ensure that DPDK applications cannot access random physical memory or memory of other process.
- Steering is done by the application when add a filter.
- A packet can be steer to multiple queue and/or application.
- Not require SR-IOV. Multiple app work over VF/PF

- Designed to be run non root thanks to verbs

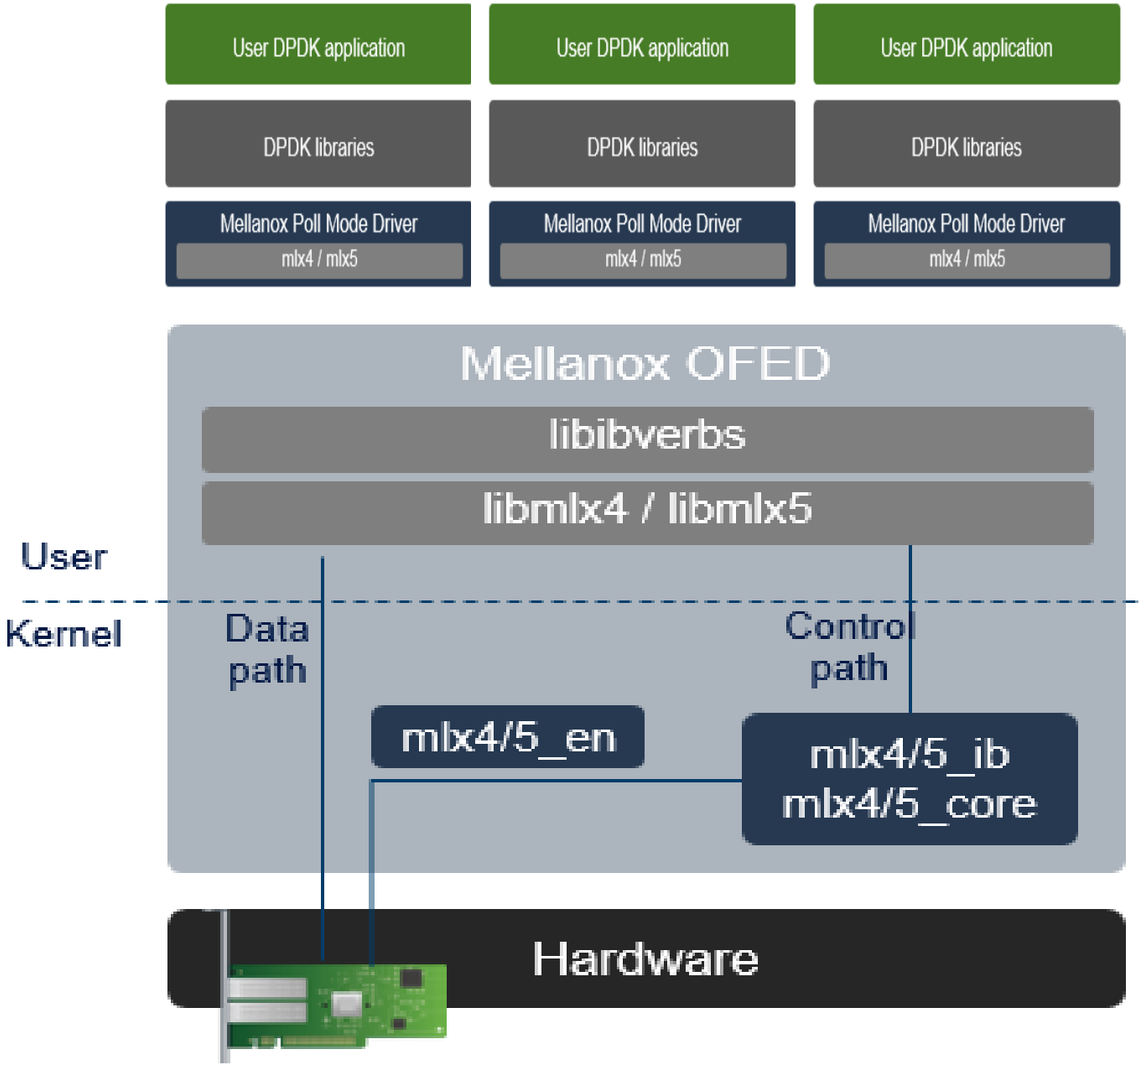

Thank You

Mellanox TECHNOLOGIES

Connect. Accelerate. Outperform.™