

Data movers are applications that

- (1) are network intensive,
- (2) process message **metadata**, but
- (3) leave message **data** unchanged.

Types of data movers

- (1) Software that processes packet **headers** but not **payload** ("NFV")
 - Examples: SW routers, NAT, load balancers, multicast collectives
- (2) Software that associate item **key** with item **data**
 - Examples: key-value stores (Memcached, ...), static webserver (Apache, ...)

Data movers waste resources

- PCIe bandwidth
- LLC space & bandwidth (DDIO)
- Memory bandwidth
- CPU cycles (if mover isn't zero-copy)

Our idea – eliminate waste with **nicmem**, the on-NIC memory

- Internal NIC SRAM used for offloading (RDMA, aRFS, SRIOV, ...)
- Typically underutilized (only 15% used in NVIDIA ConnectX NICs by default)
- Can easily/cheaply increase its size (~ \$0.2 per MB)

Specifics

- (1) **Expose** unused nicmem to applications, as regular memory, via MMIO
- (2) **Split** headers & data on receive – headers to hostmem & payload to nicmem
- (3) **Inline** headers – use only one descriptor per-packet

