



THE INFINIBAND / RDMA PANEL

# WHAT ????

- Infiniband is a high speed network (40Gb, 56Gb, 100Gb).
- Direct Host to Host, Memory to memory transfers (RDMA)
- Kernel bypass and Offload API (Verbs/RDMA API)



# KERNEL OFFLOAD

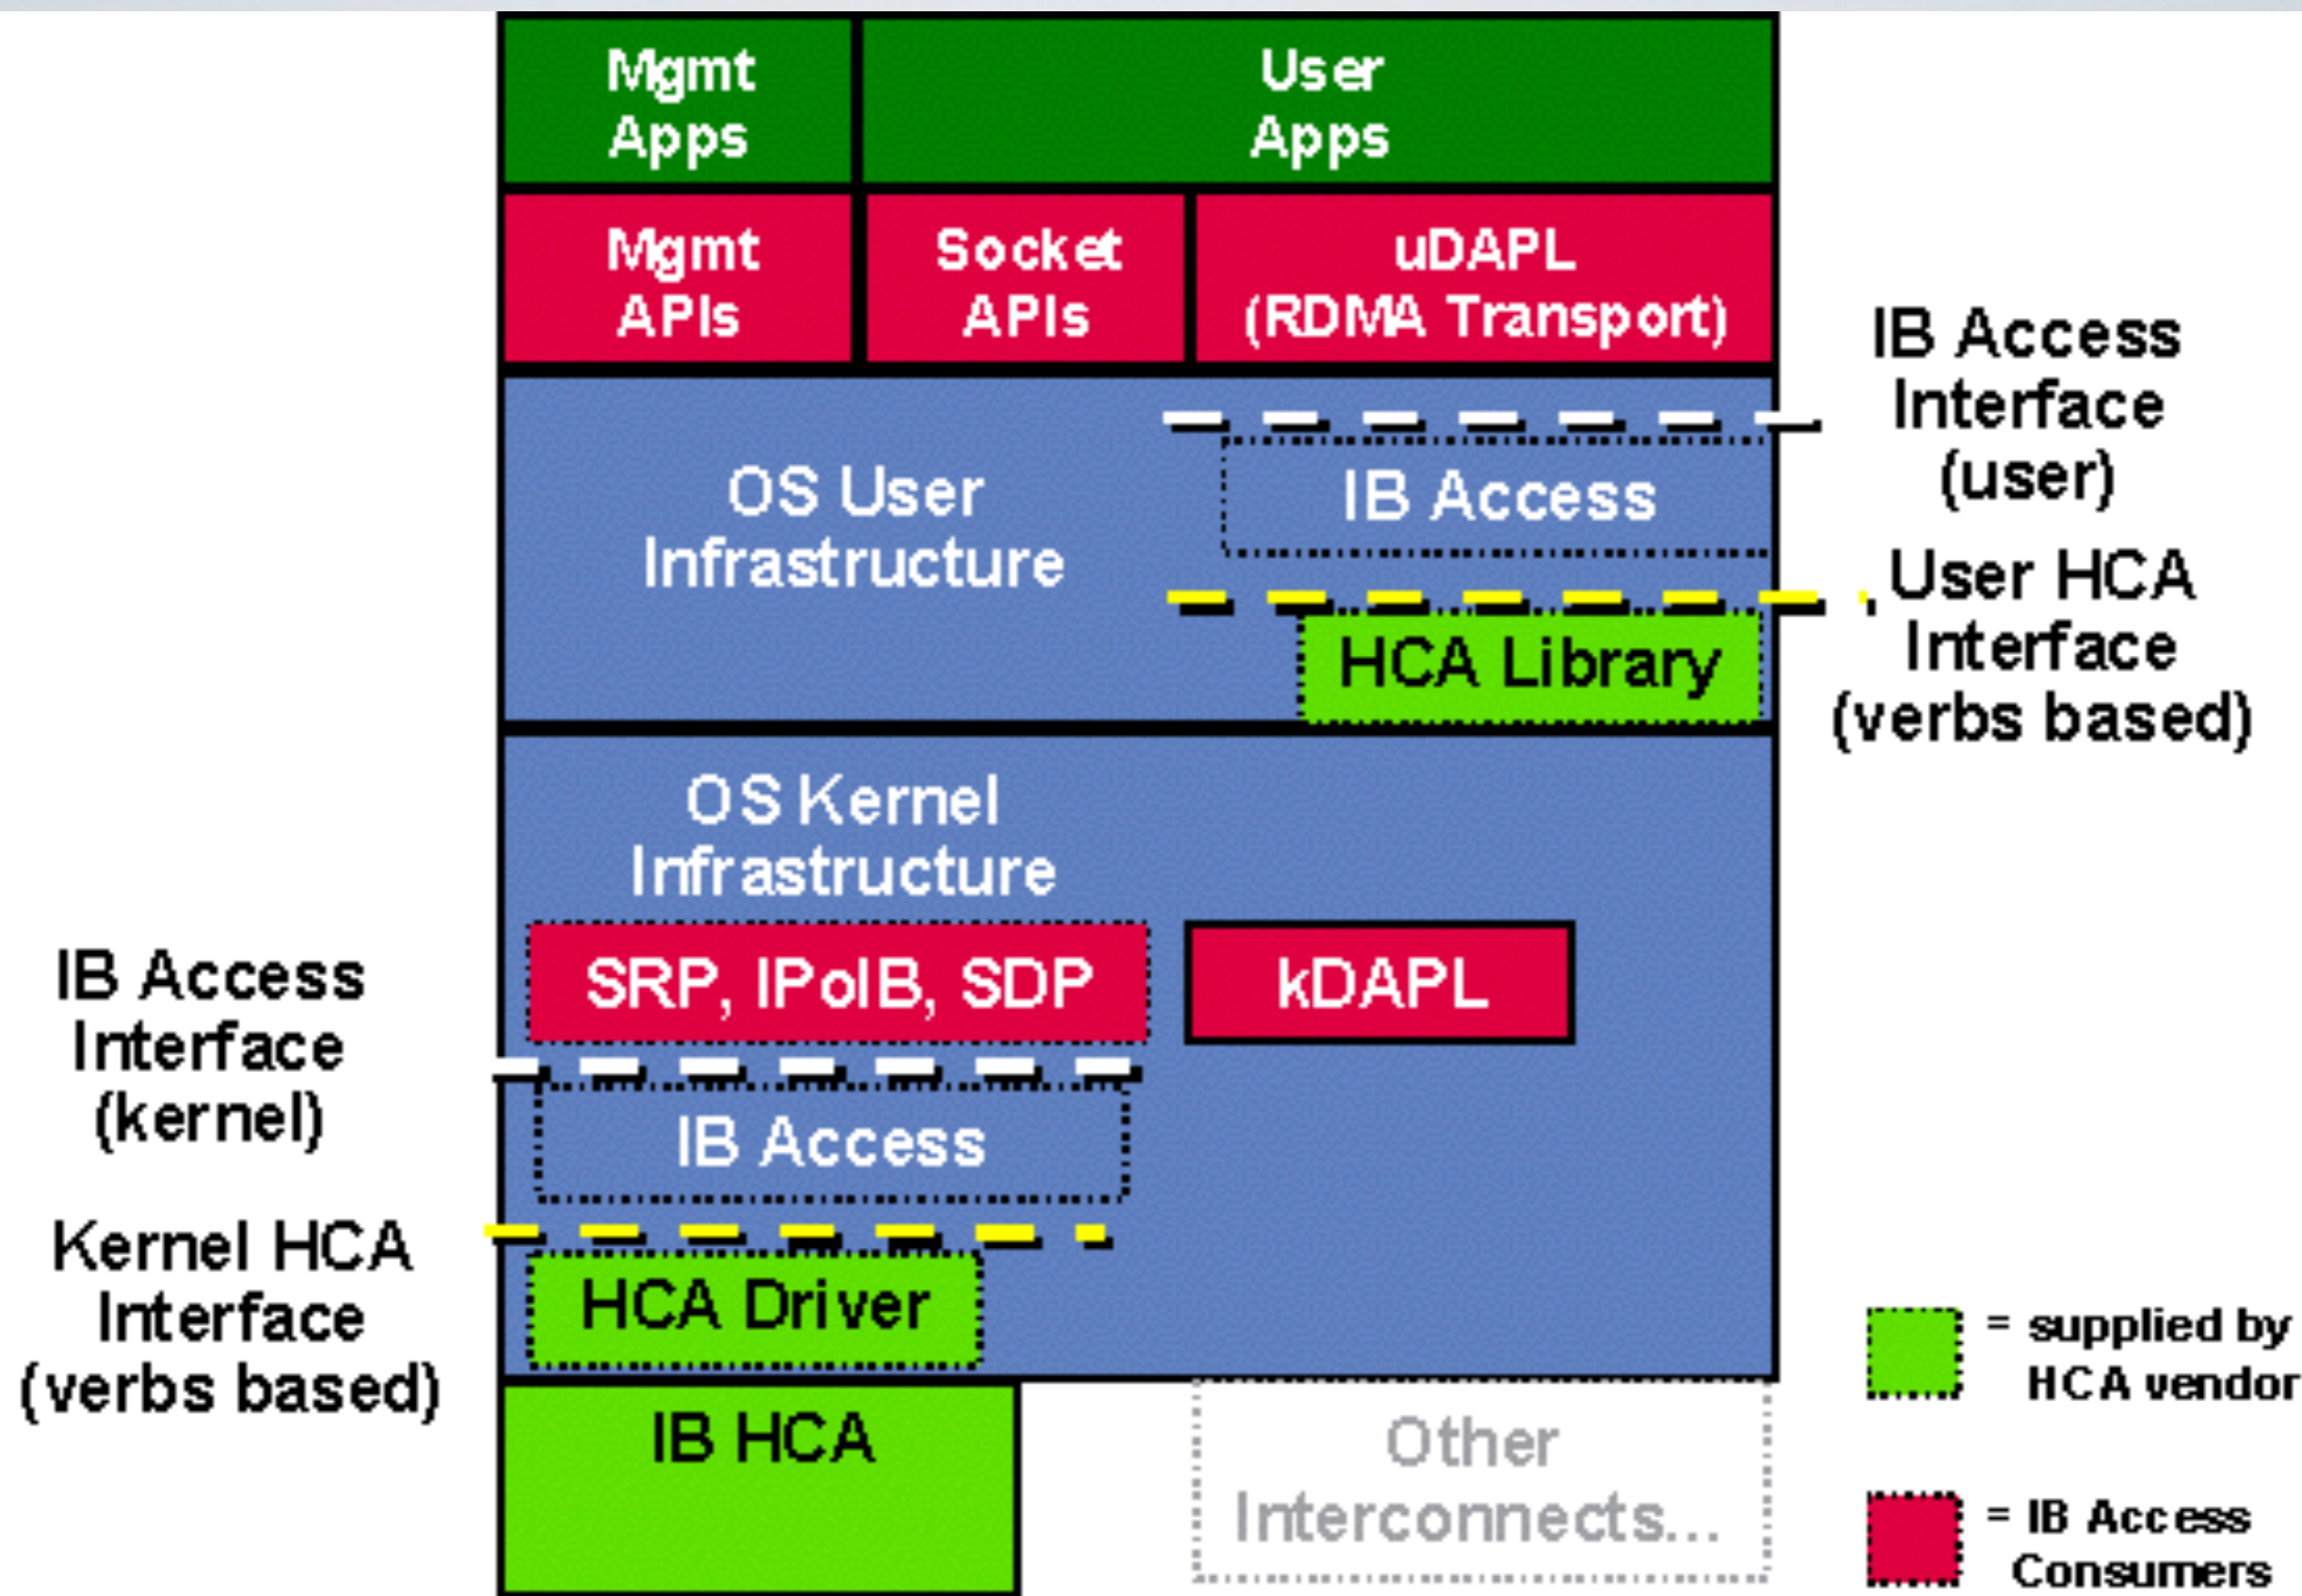
- Kernel provides API to setup and configure network transfers.
- Kernel is not involved in actual packet receive and send.
- Registration of memory regions for RDMA
- One sided transfers (RDMA proper).
- Messaging (TX, RX rings in user space)

# PROTOCOLS SUPPORTED BY RDMA


- Infiniband (SDR,DDR,QDR, FDR, EDR) (10-100Gb)
- ROCE (Full RDMA API via Ethernet)
- iWARP (One sided RDMA via TCP/IP)
- Raw Ethernet (RDMA Messaging API)
- PCI-E (Xeon Phi)
- Future: SSD, memory mapped devices.

# PANEL

- Jason Gunthorpe (Obsidian)
- Or Gerlitz (Mellanox)
- Sean Hefty (Intel)
- Roland Dreier (Pure Storage, Maintainer of the RDMA kernel subsystem)



Per 1x Lane Bandwidth (Gb/s)			
DDR	QDR	FDR	EDR
5Gb/s	10Gb/s	14Gb/s	26Gb/s



### Traditional Model

### Kernel Bypass Model

