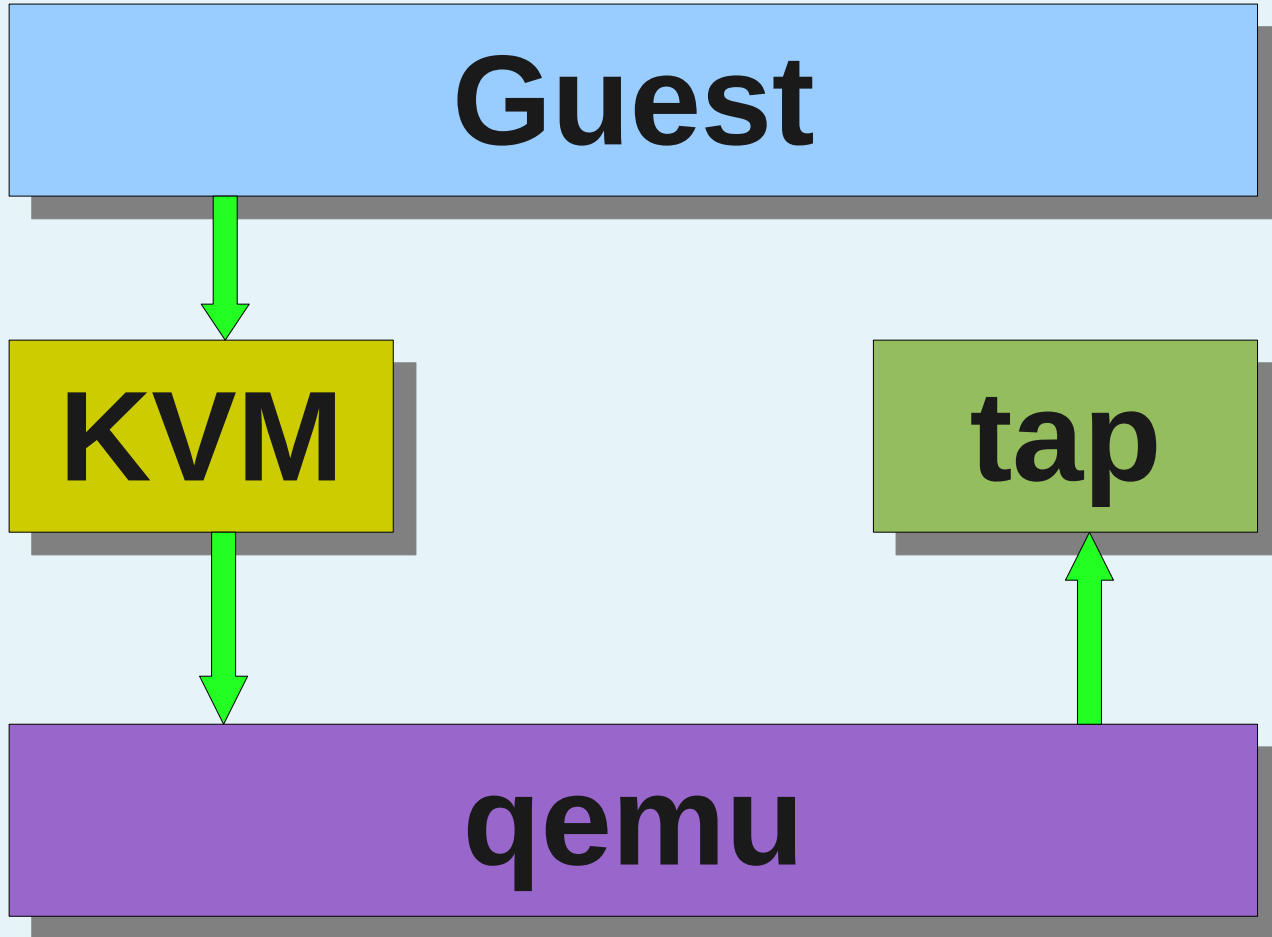


virtio- and vhost- net


Michael S Tsirkin Red Hat

need for speed
Performance challenges

TX



RX



Vhost-net

Guest

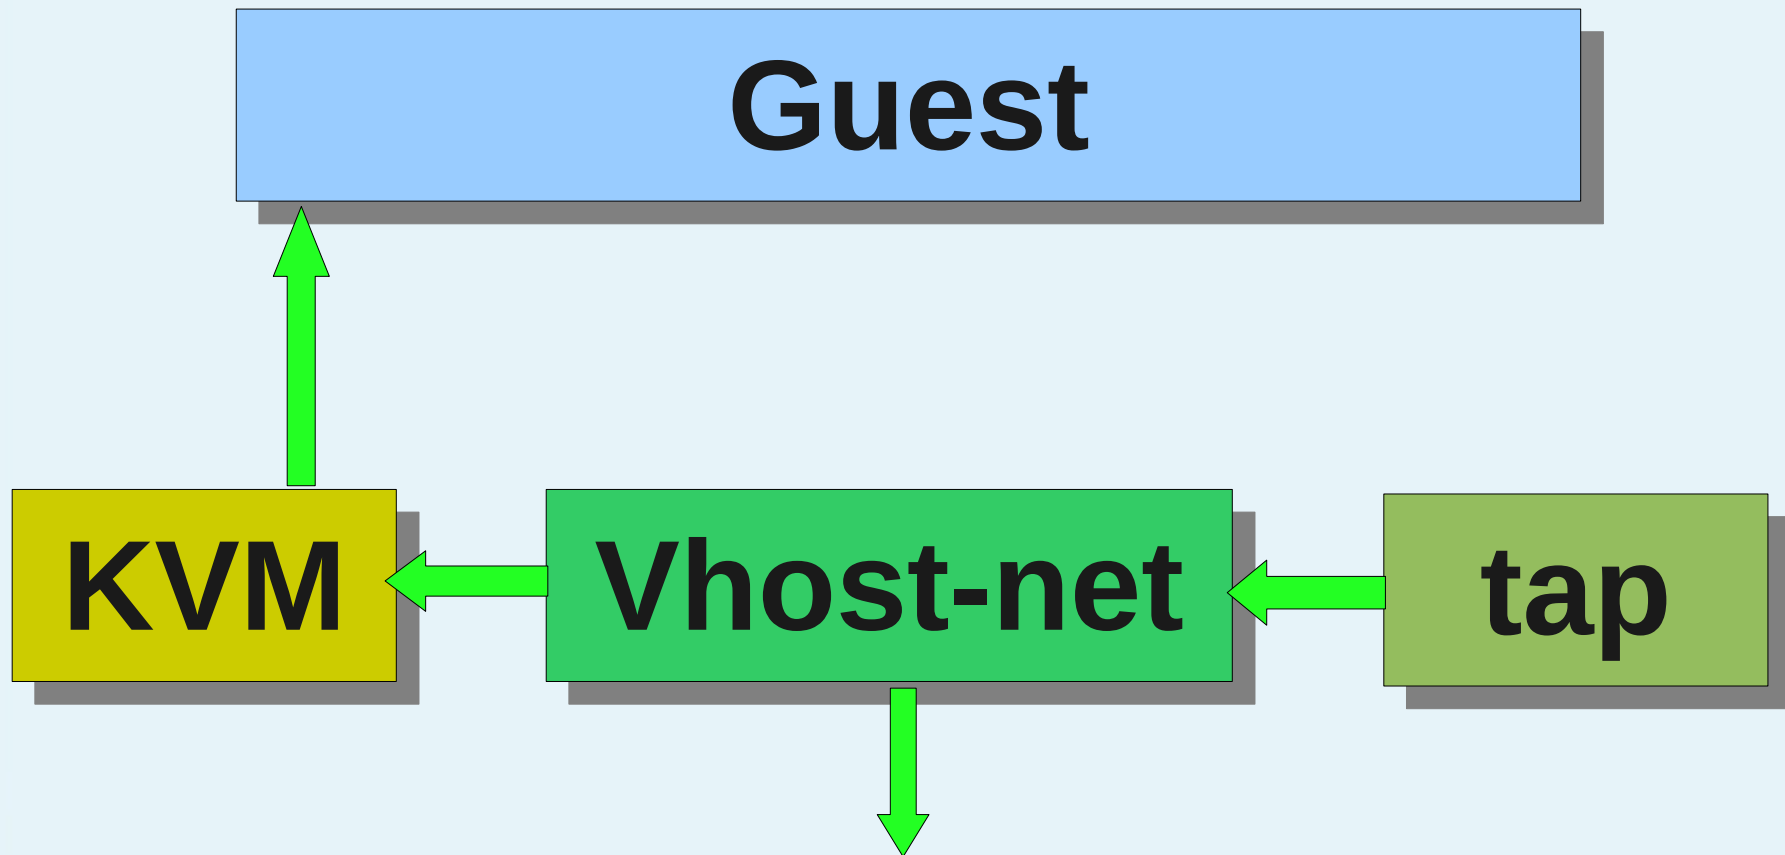
KVM Vhost-net tap

qemu

TX



RX



ioeventfd

Guest

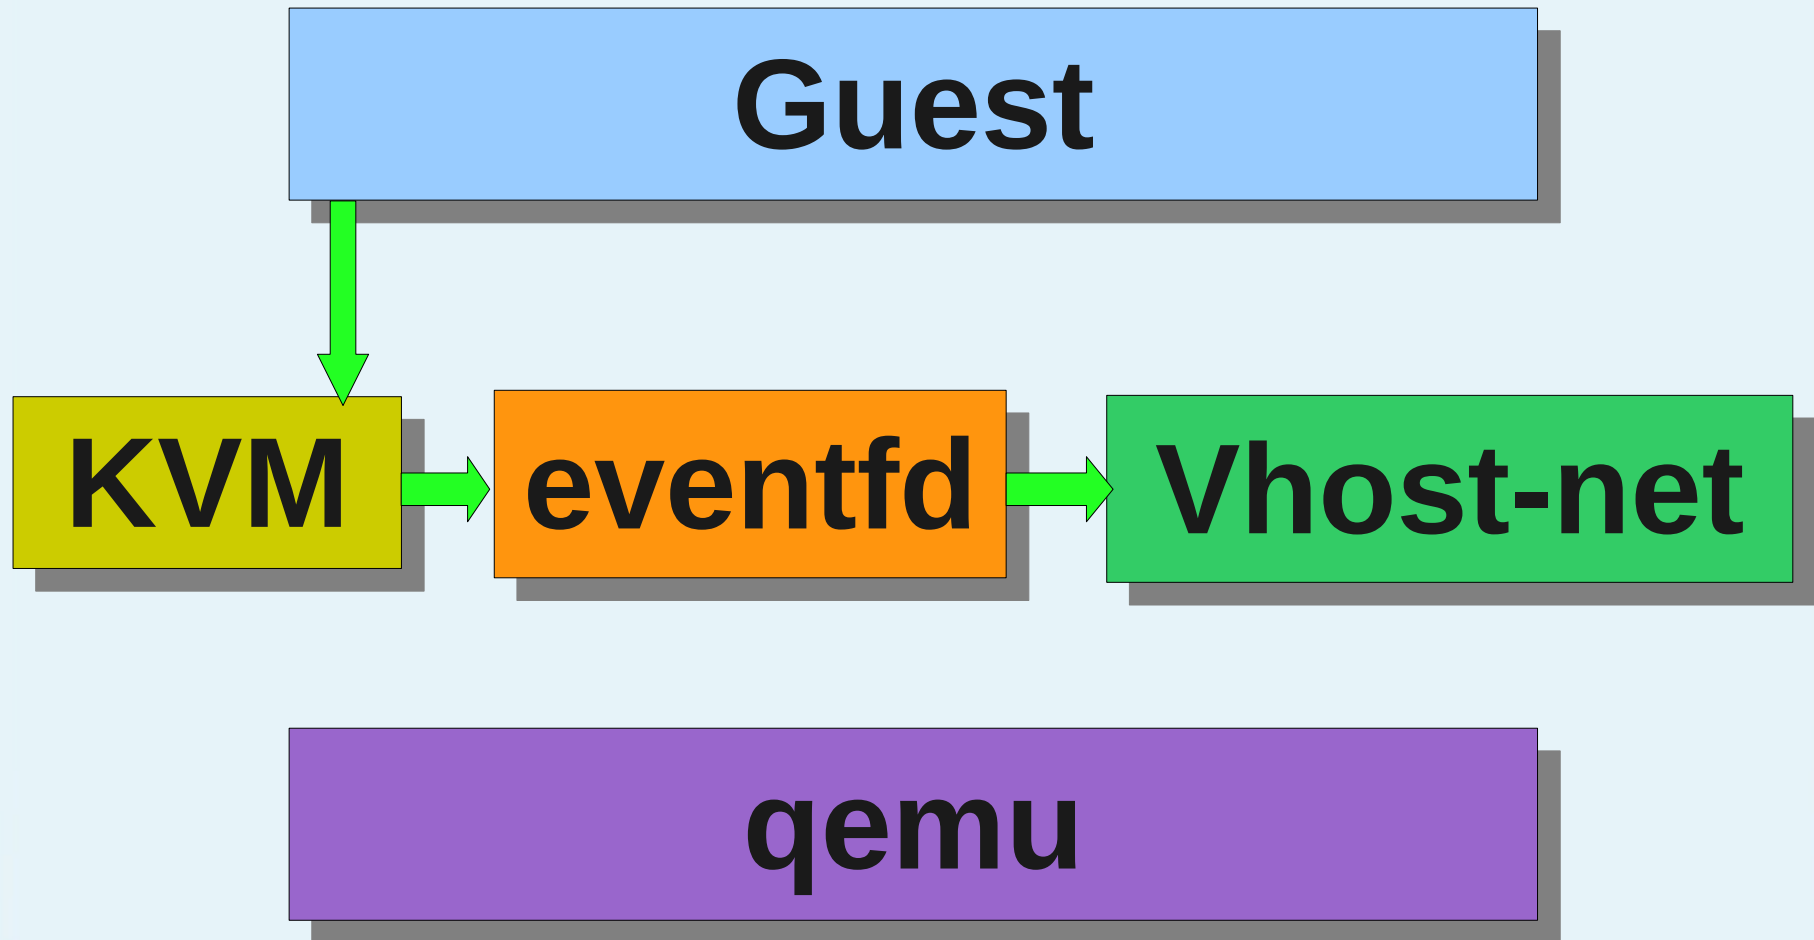
KVM

eventfd

Vhost-net

qemu

ioeventfd



irqfd

Guest

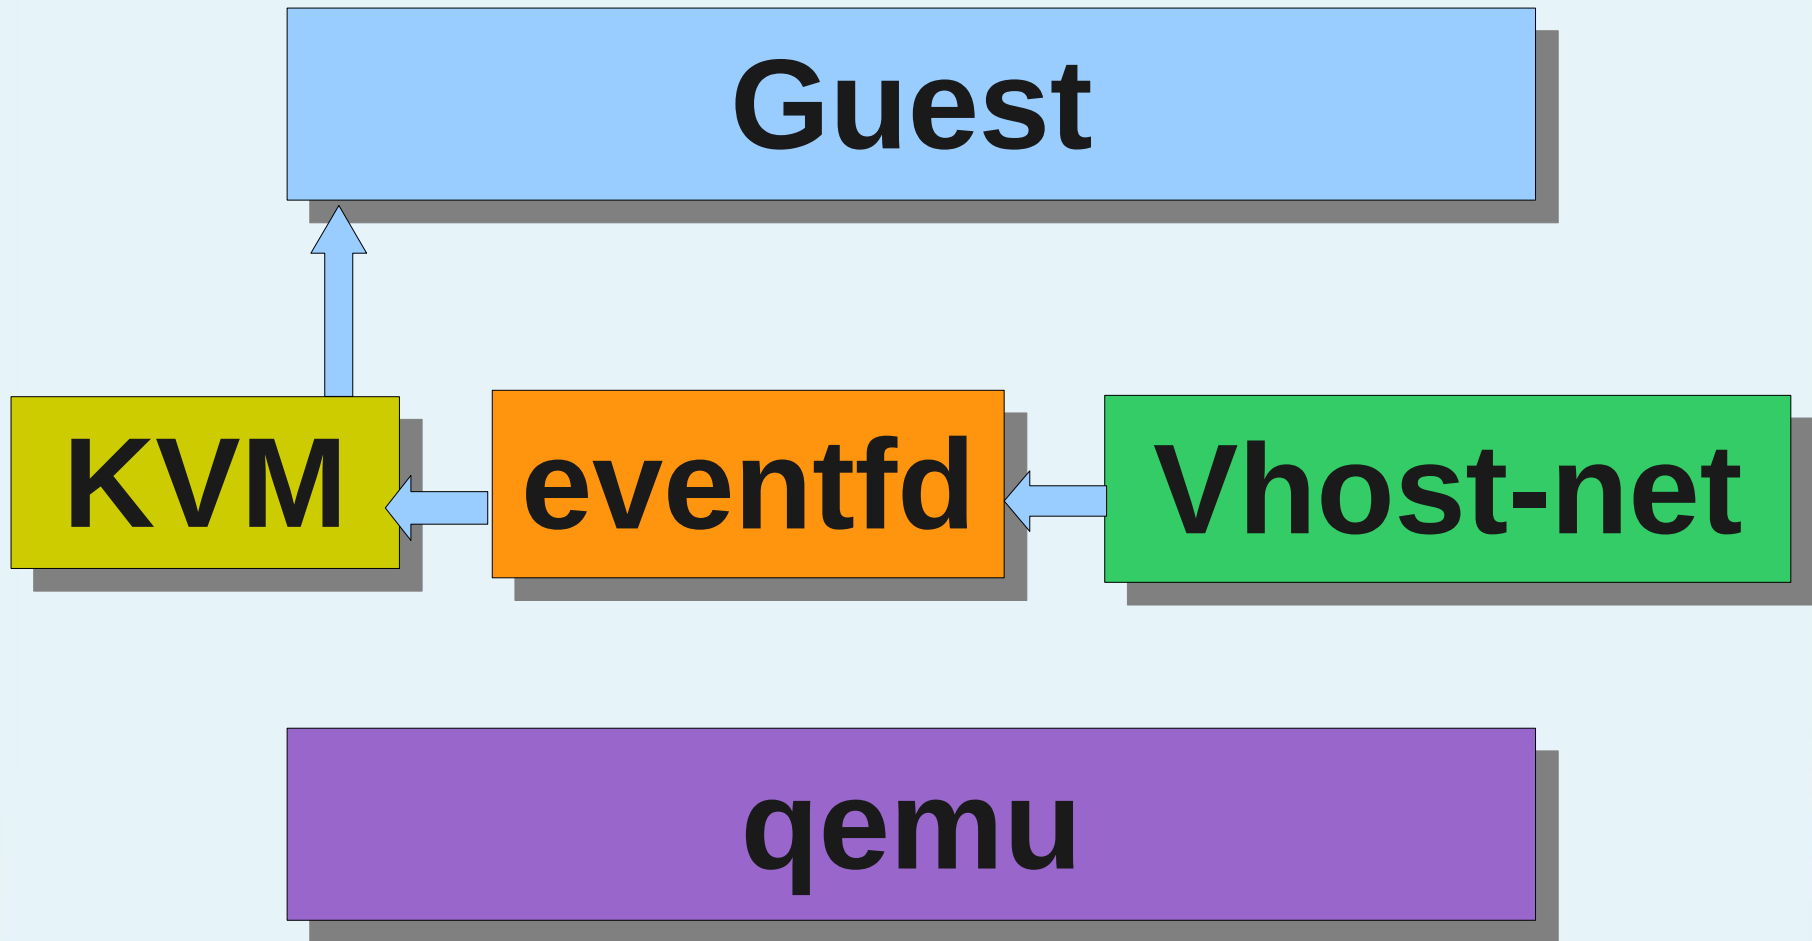
KVM

eventfd

Vhost-net

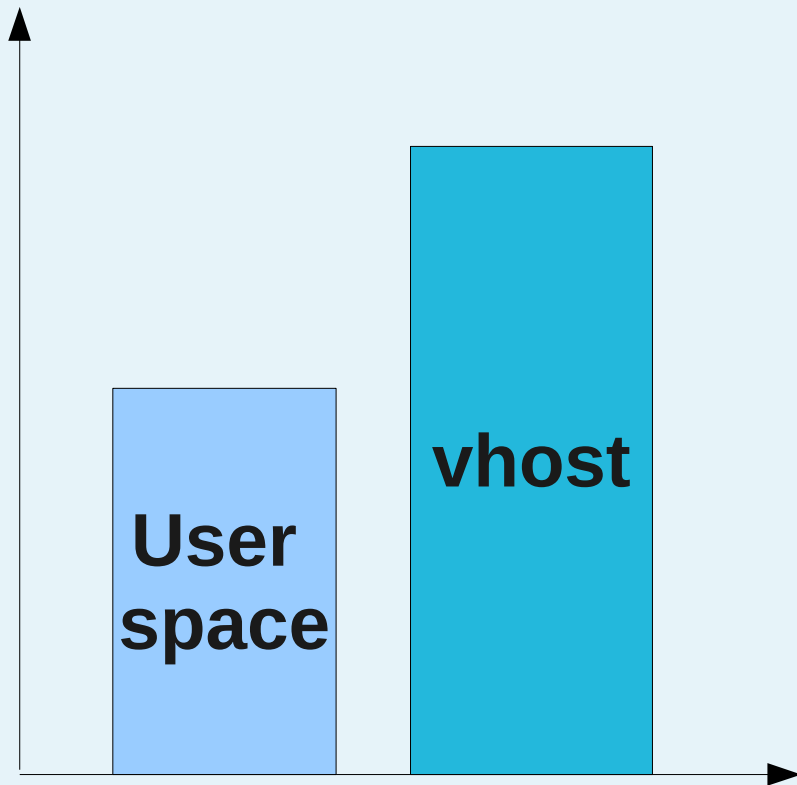
qemu

irqfd

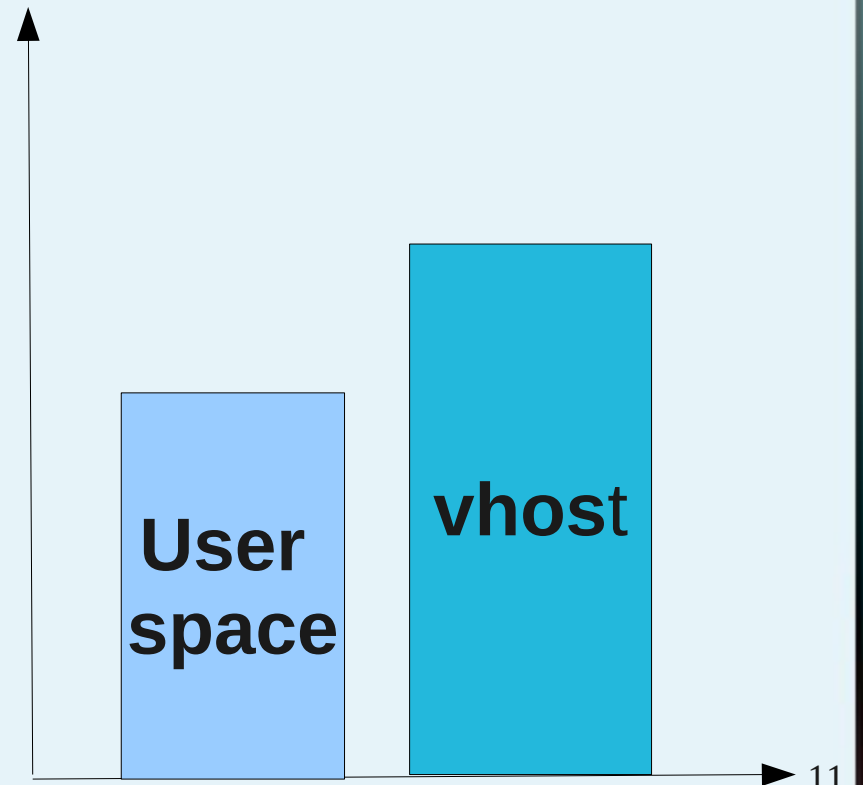


Mb/s

- Guest → Host



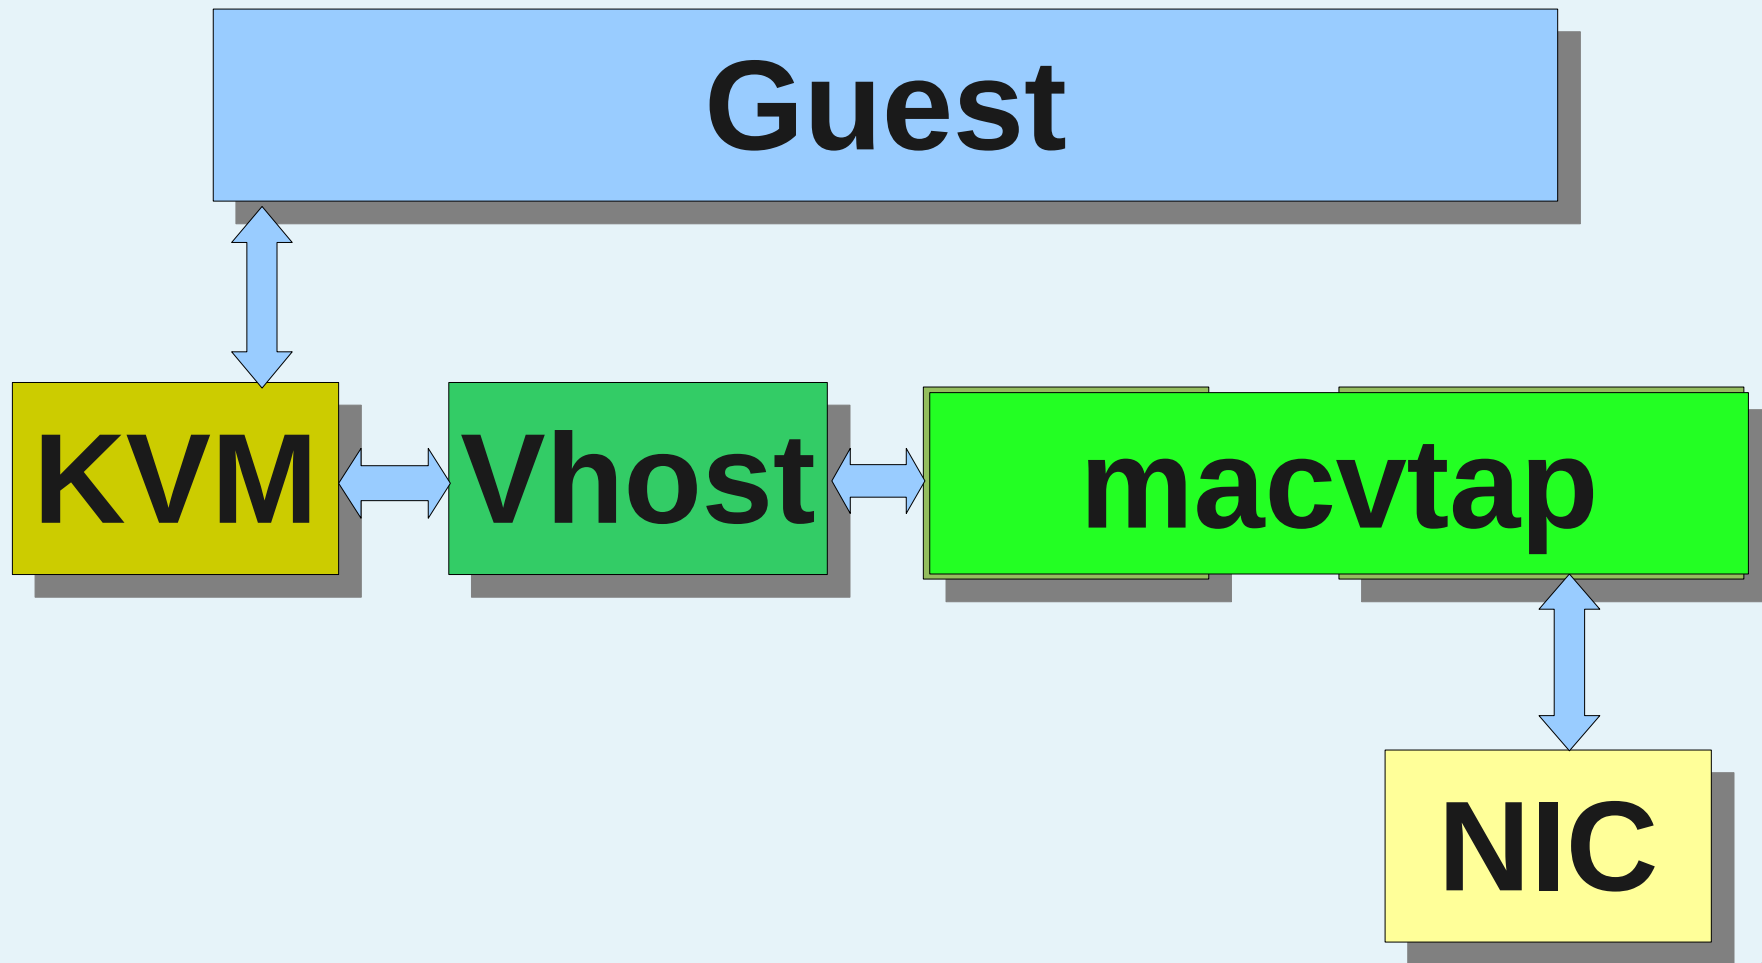
- Host → Guest



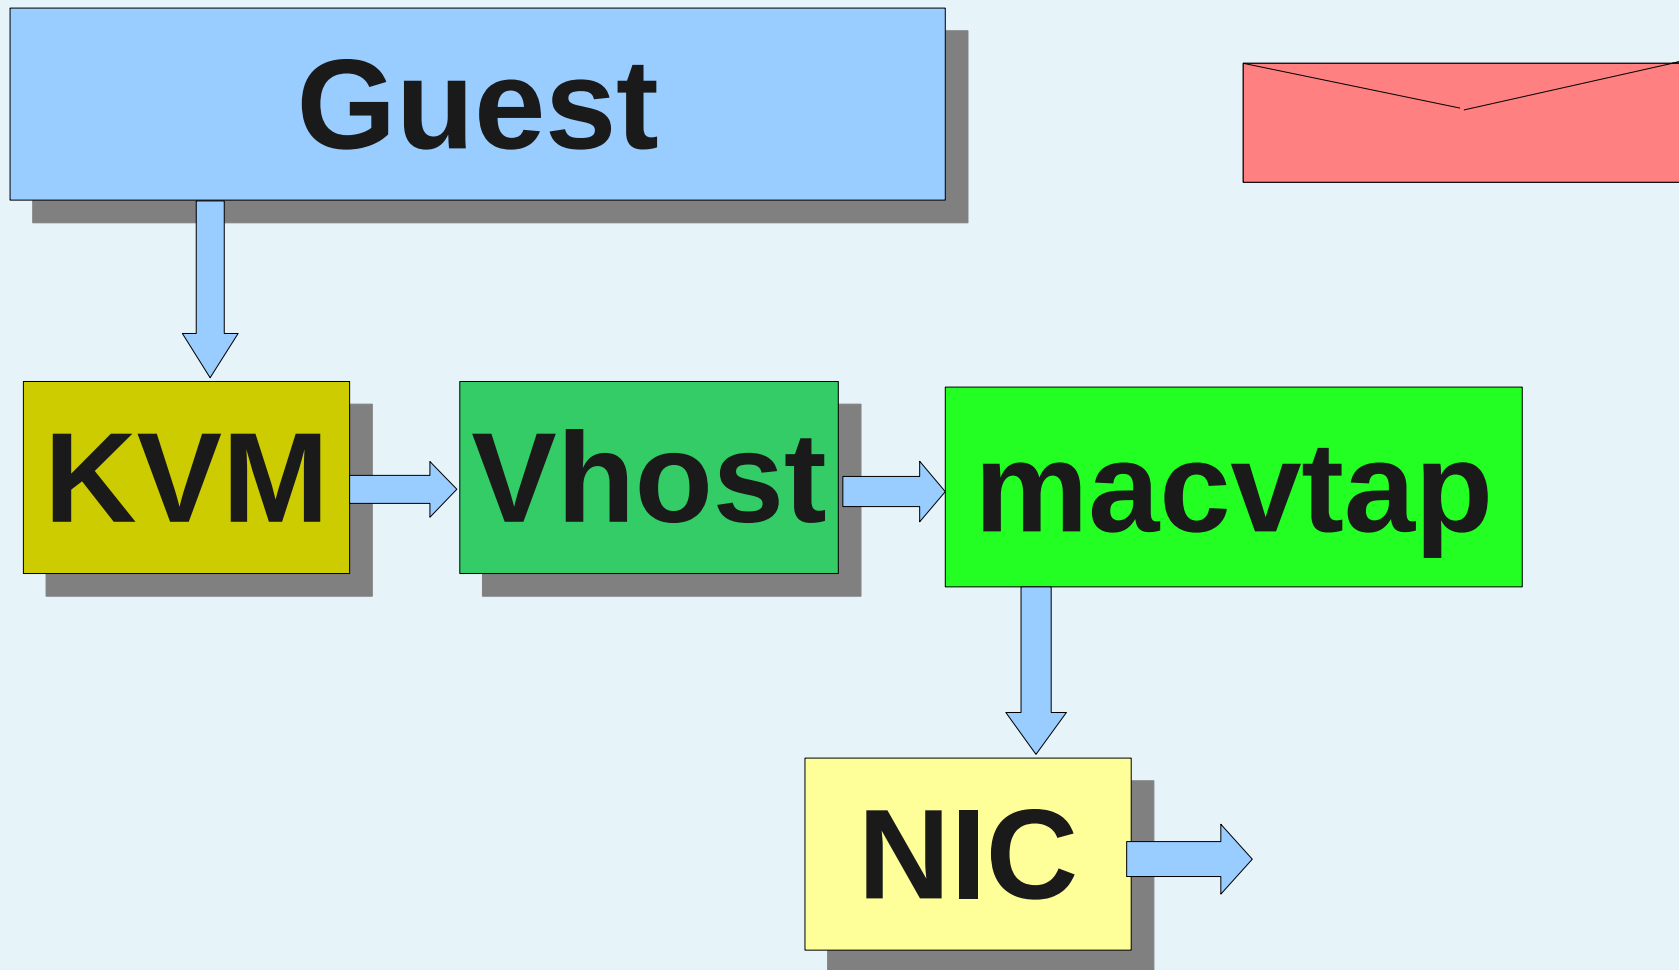
trans/s



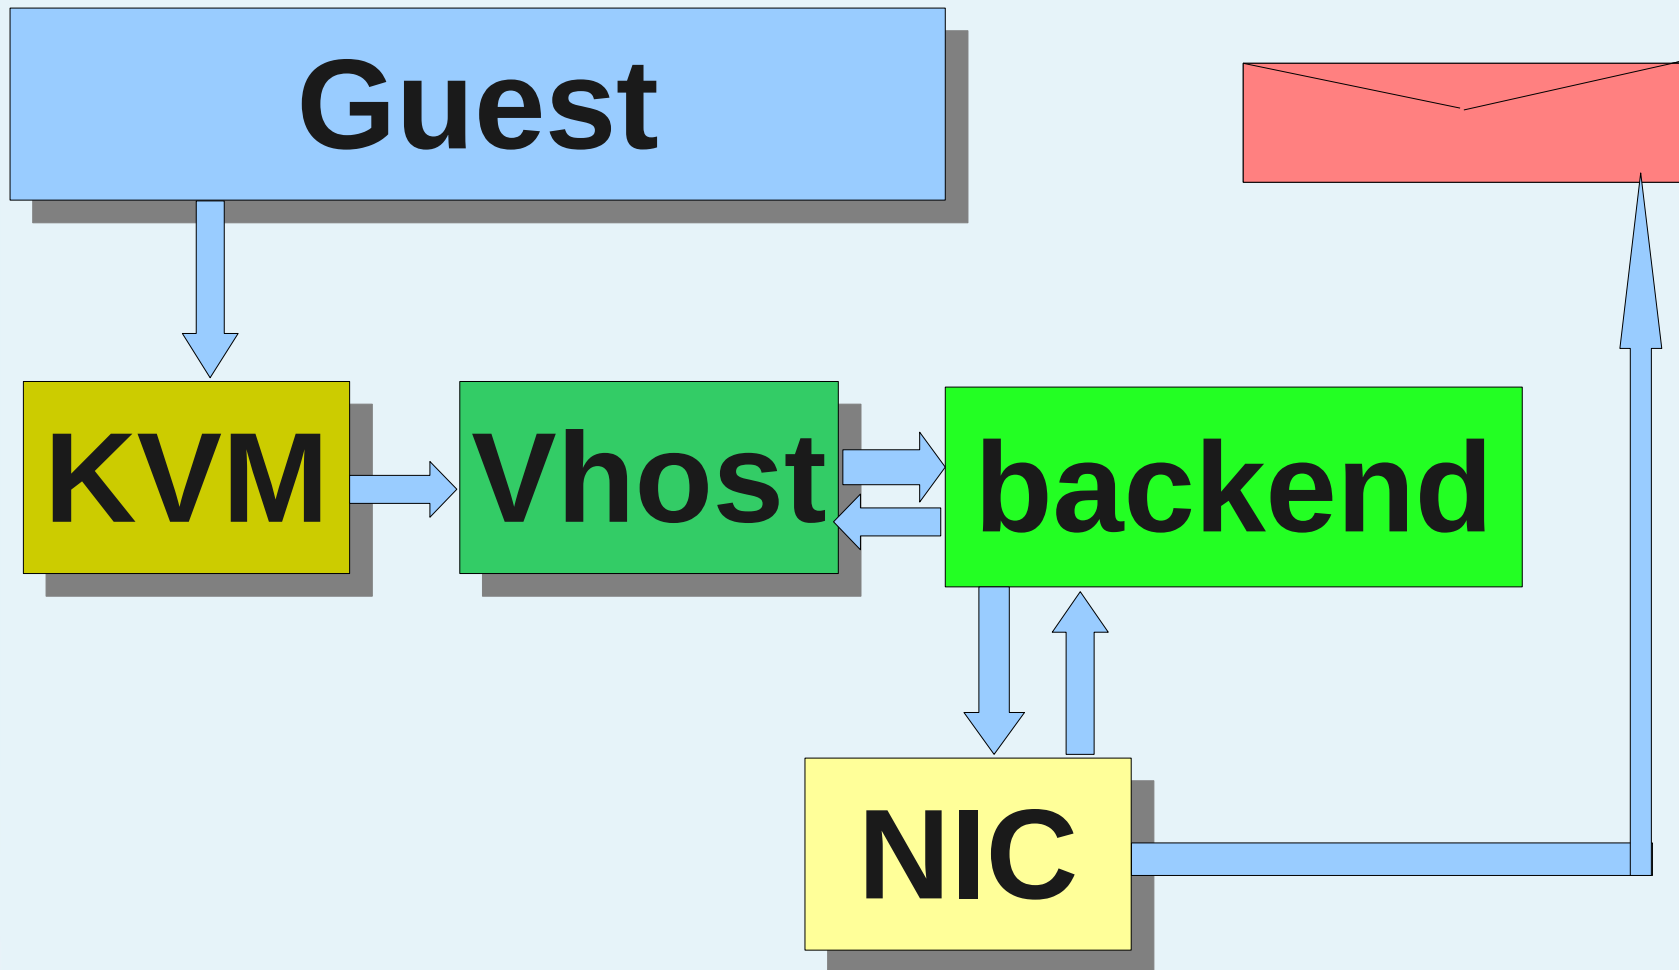
macvtap



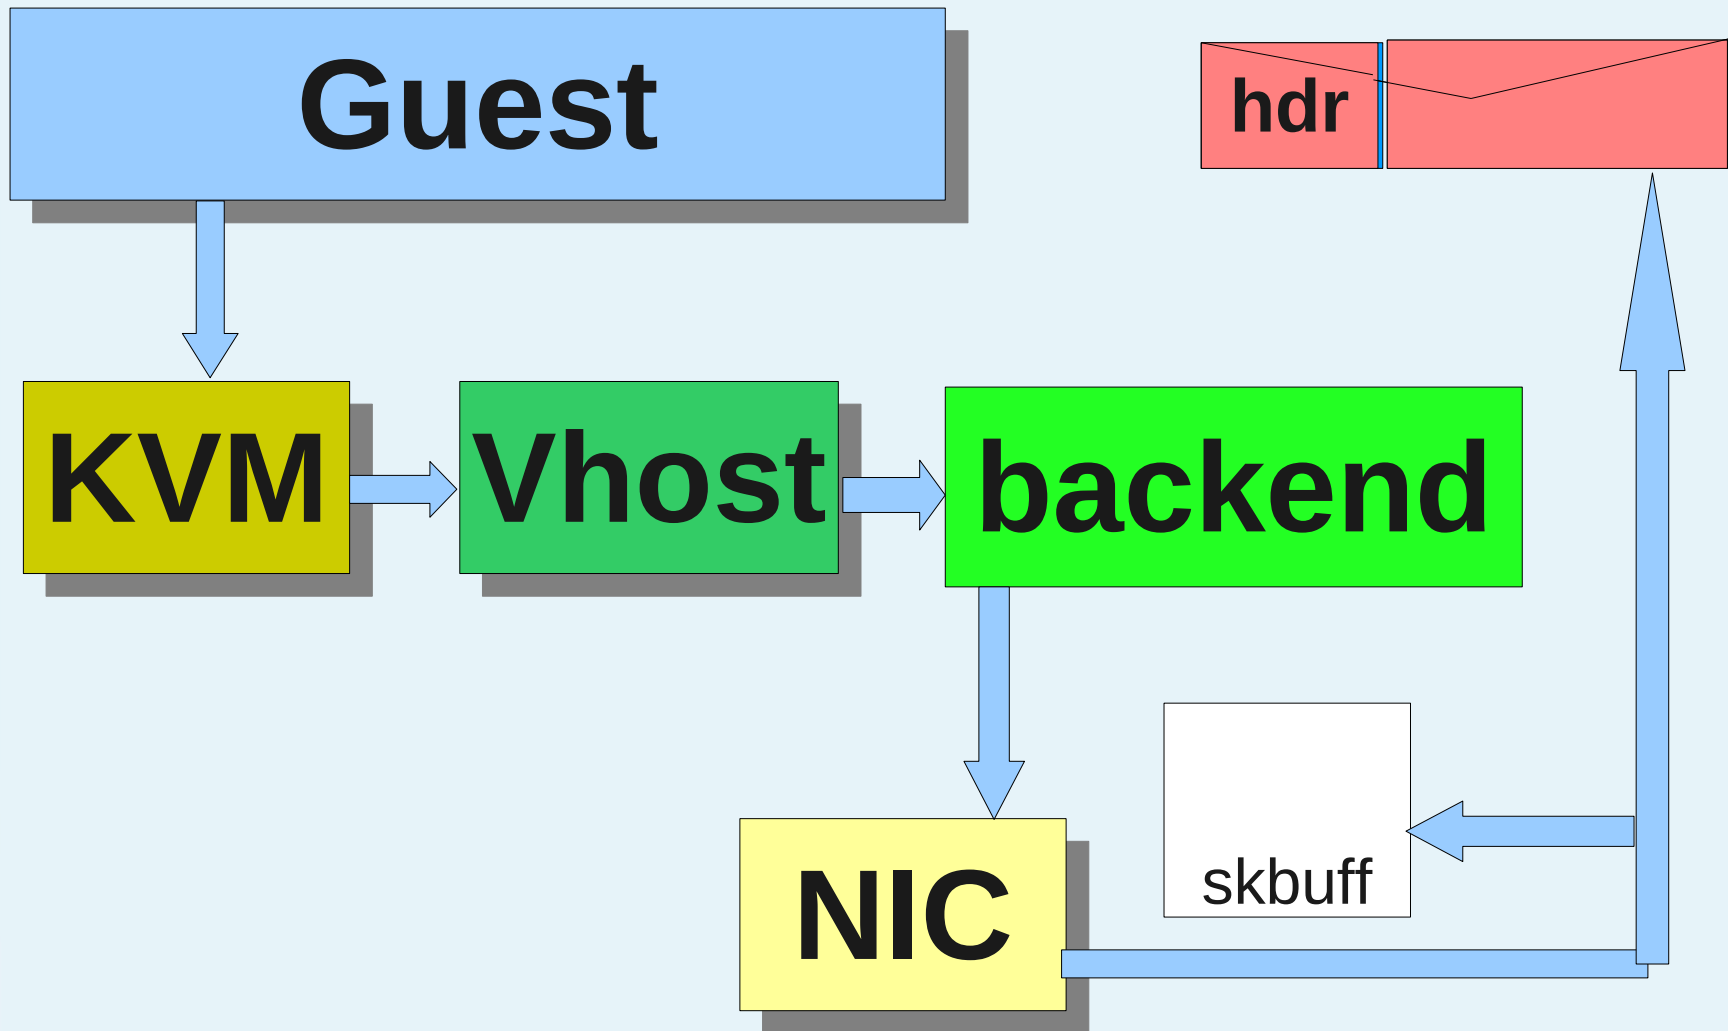
zcopy



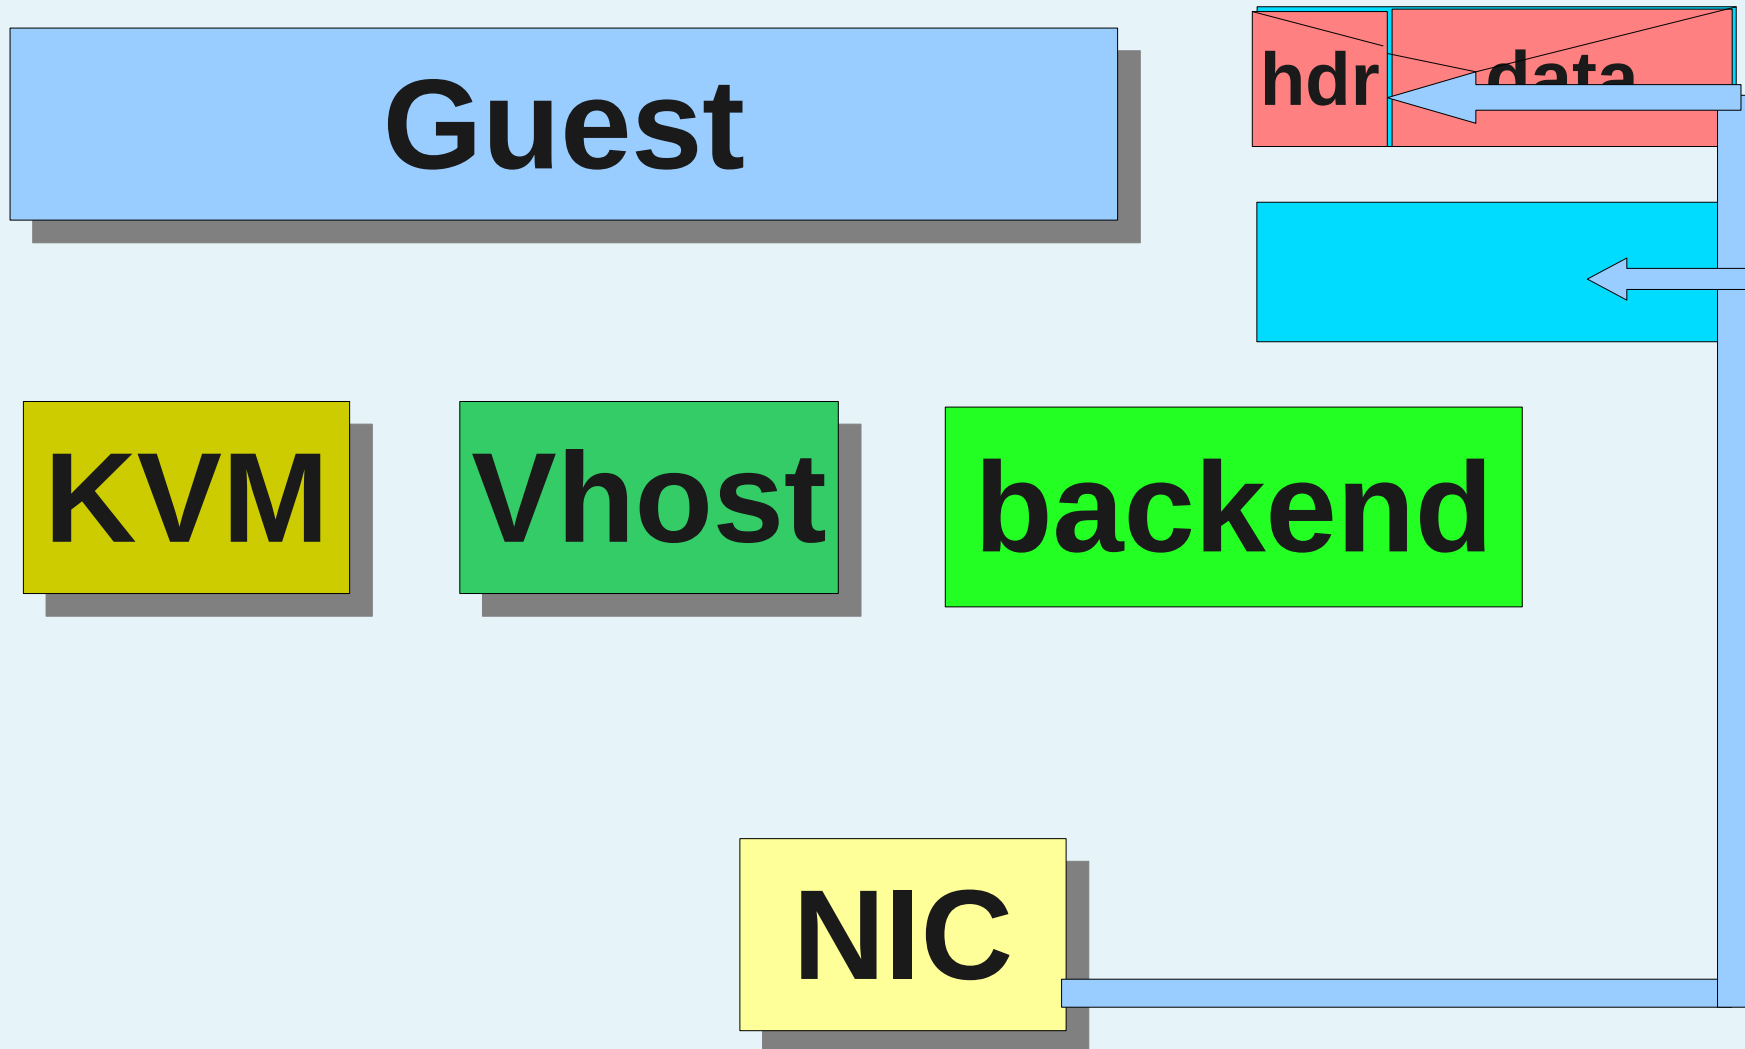
zcopy



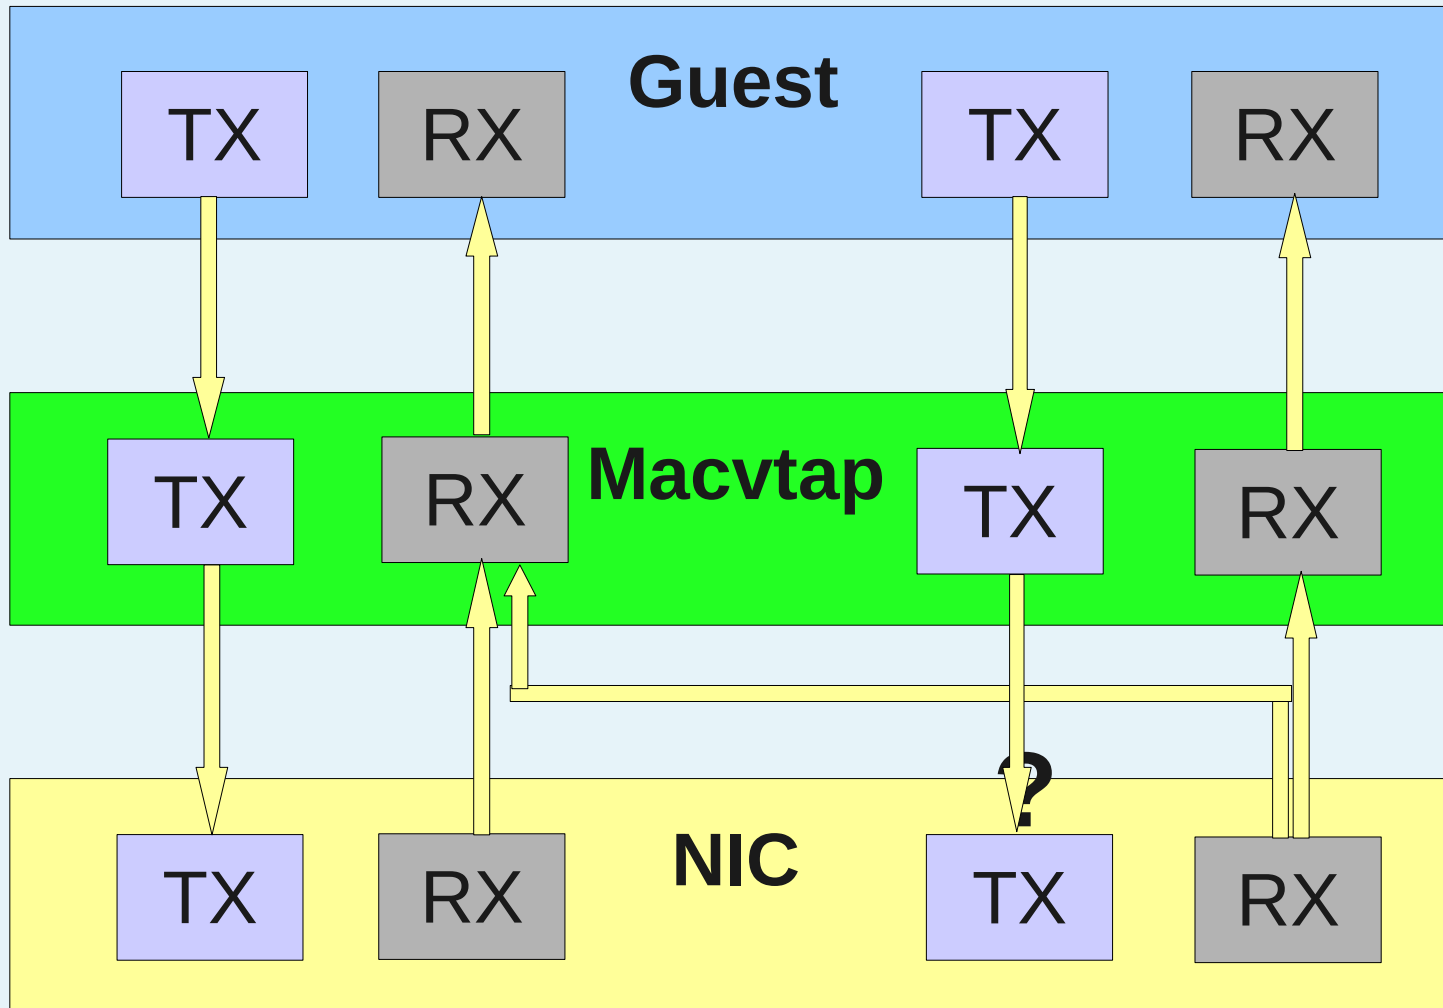
Header copy



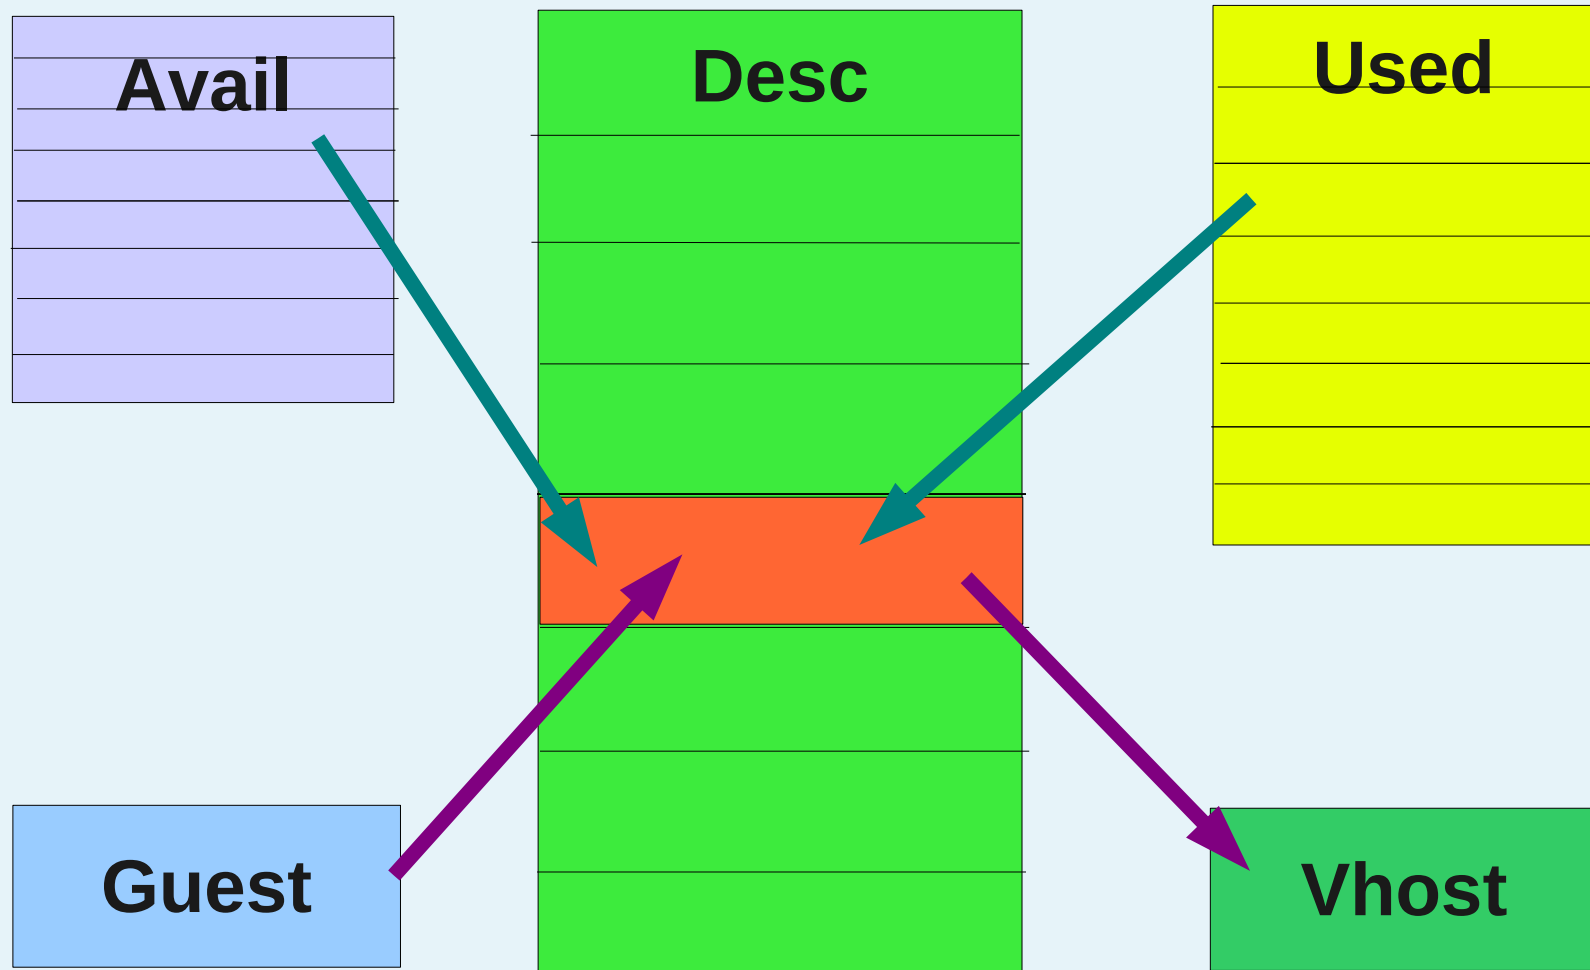
alignment



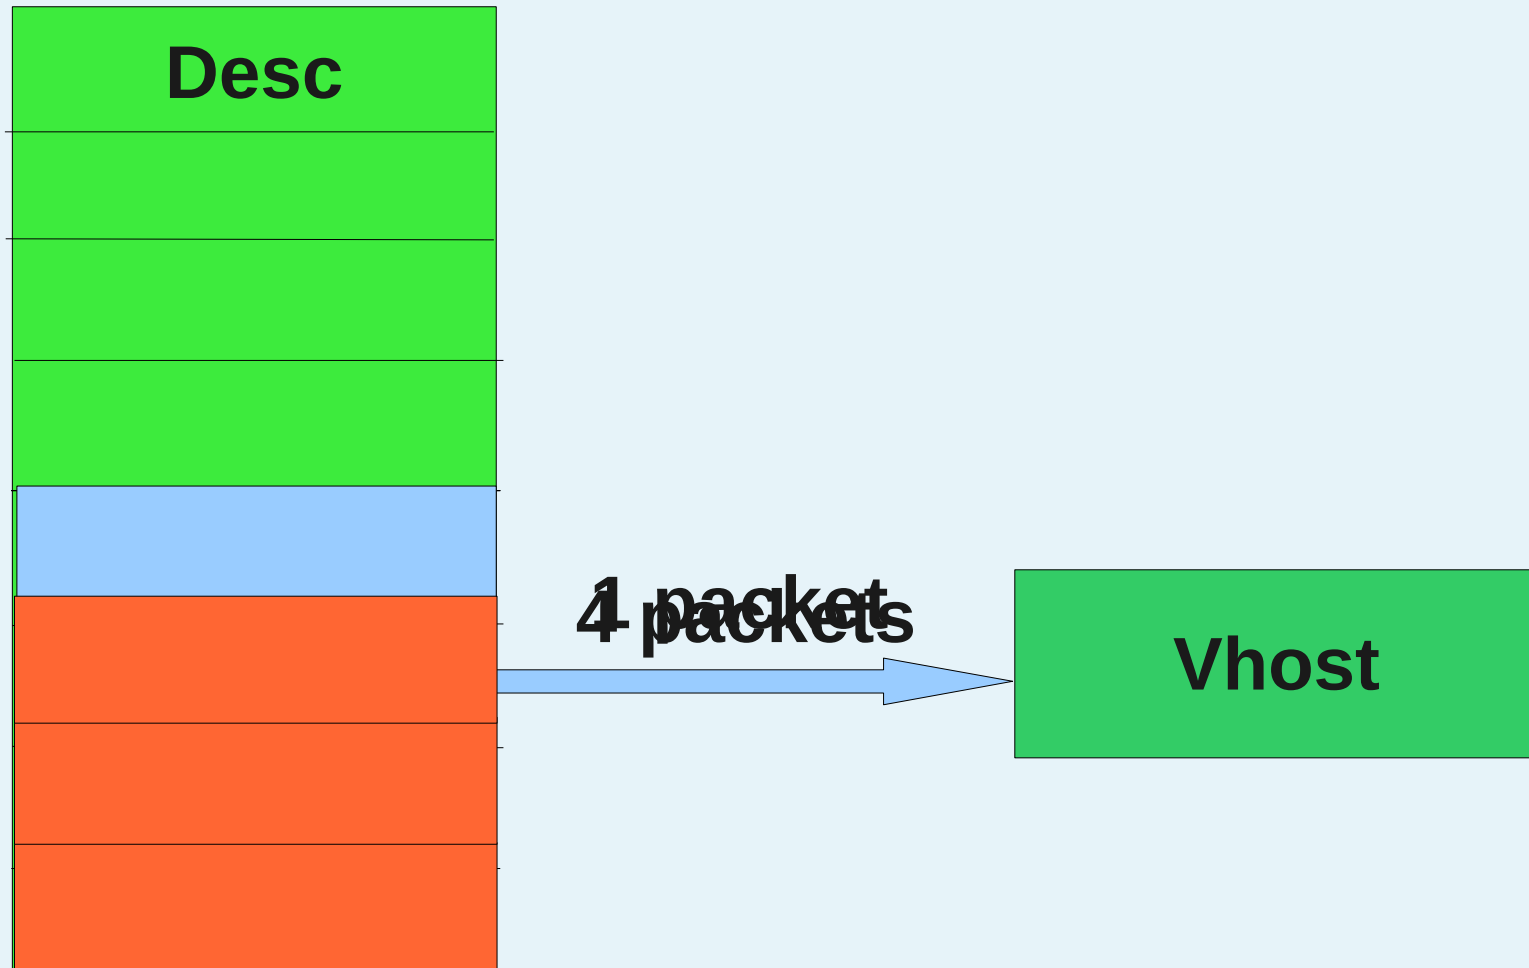
multiqueue



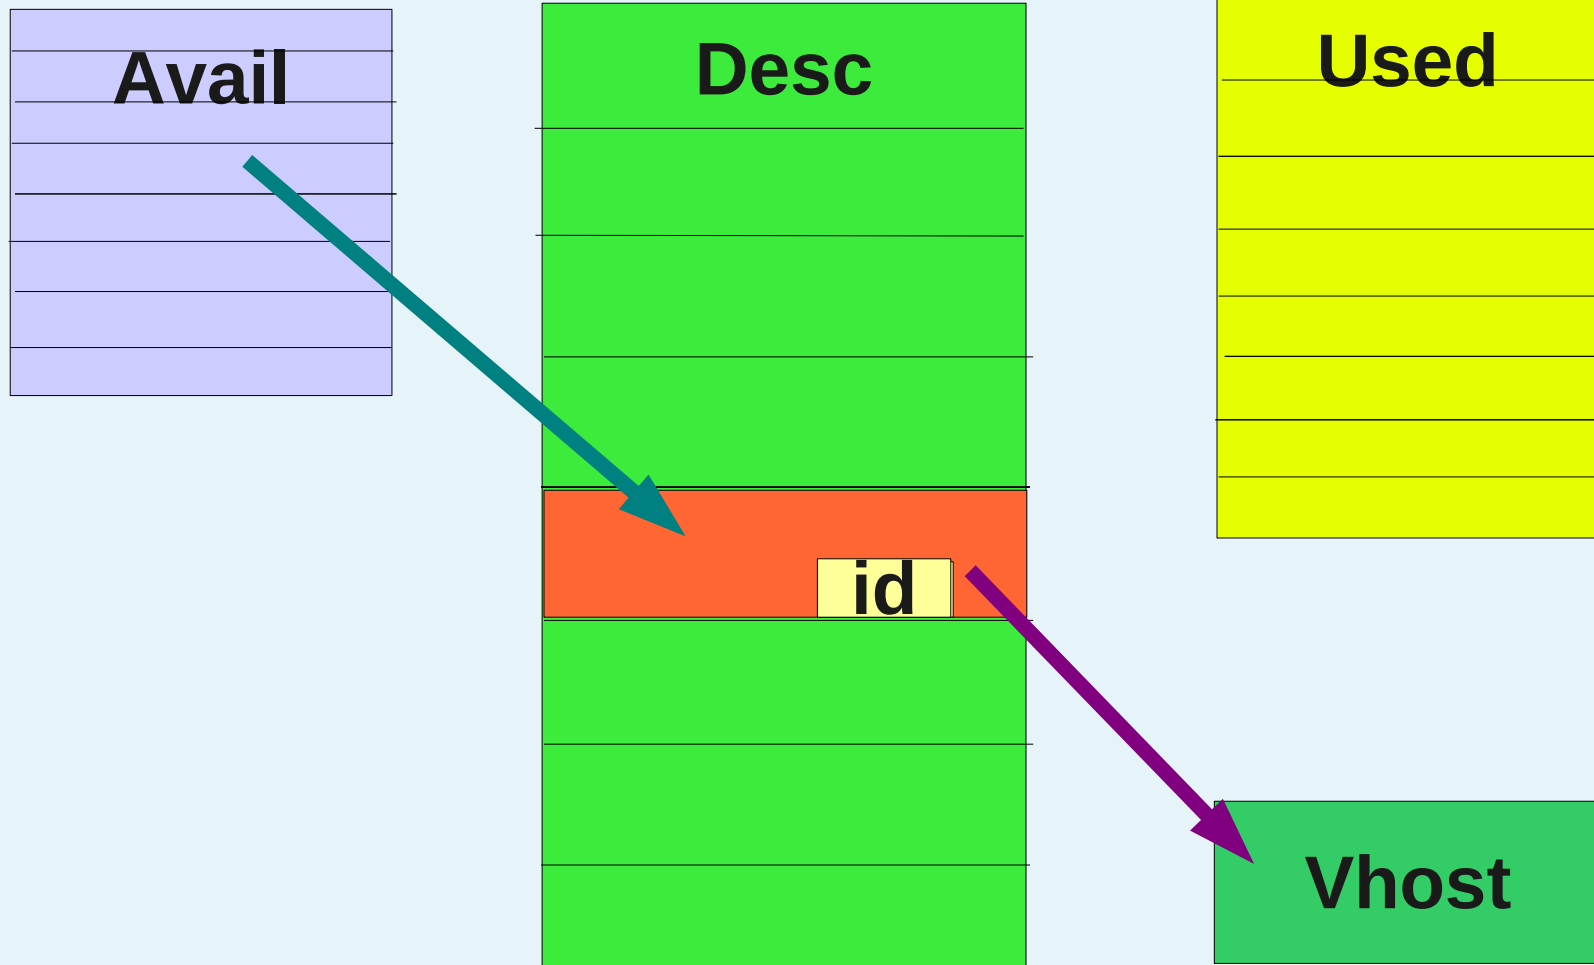
Virtio ring




fragmentation



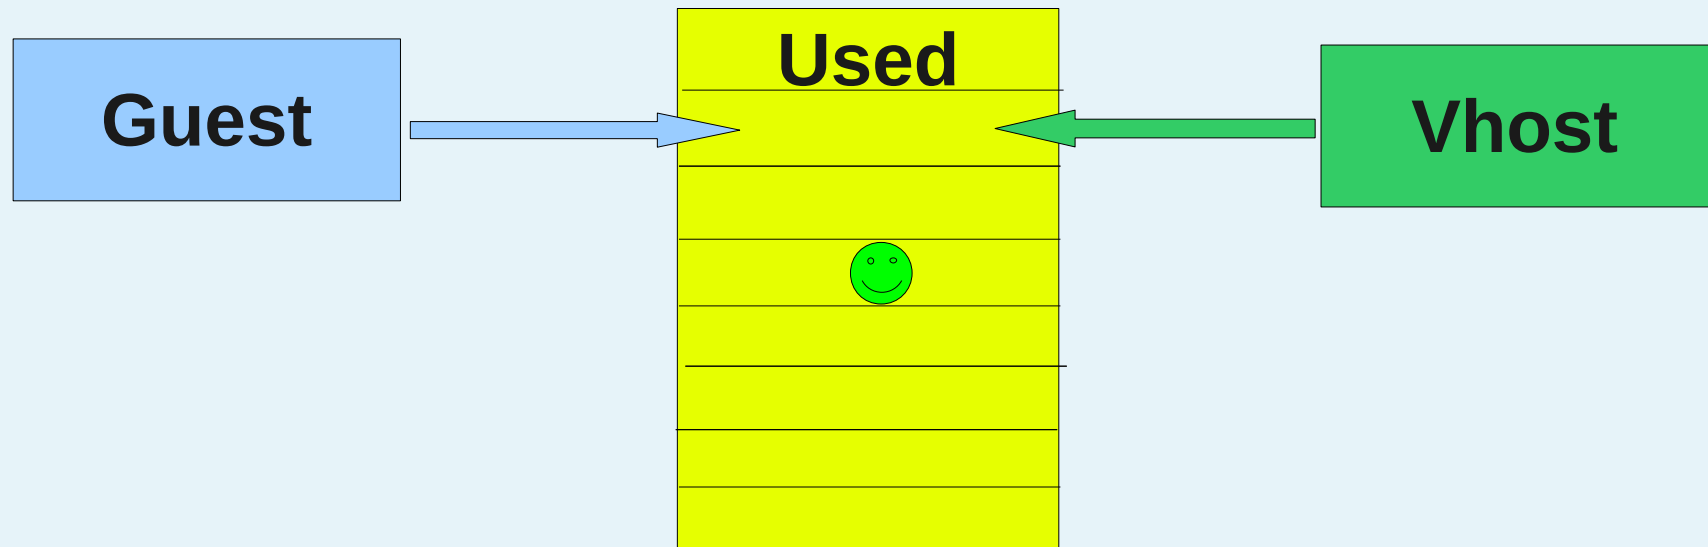
new ring




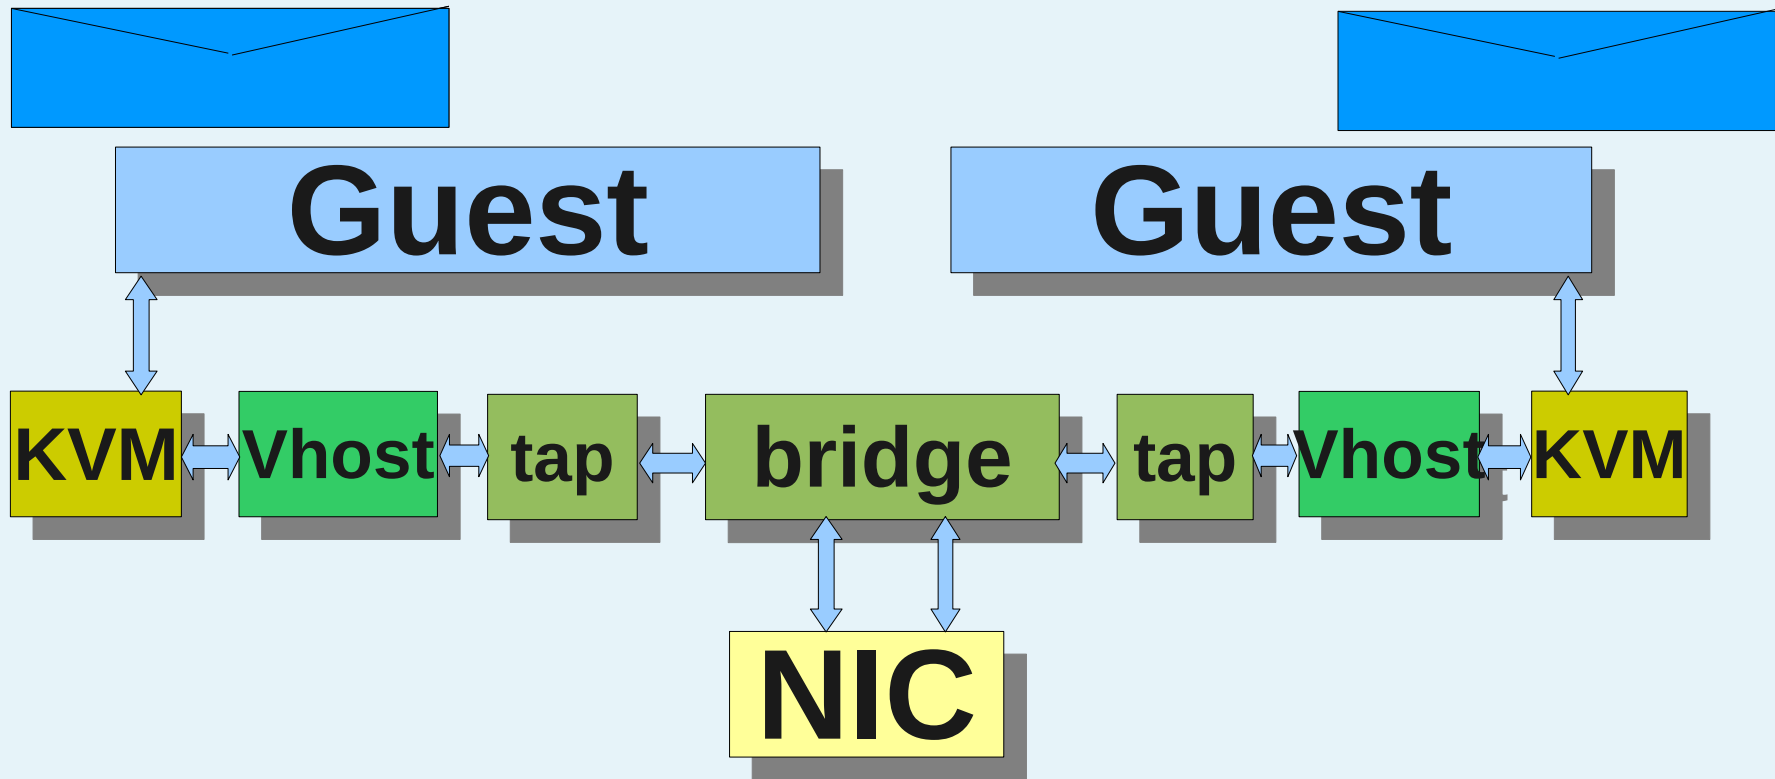
window




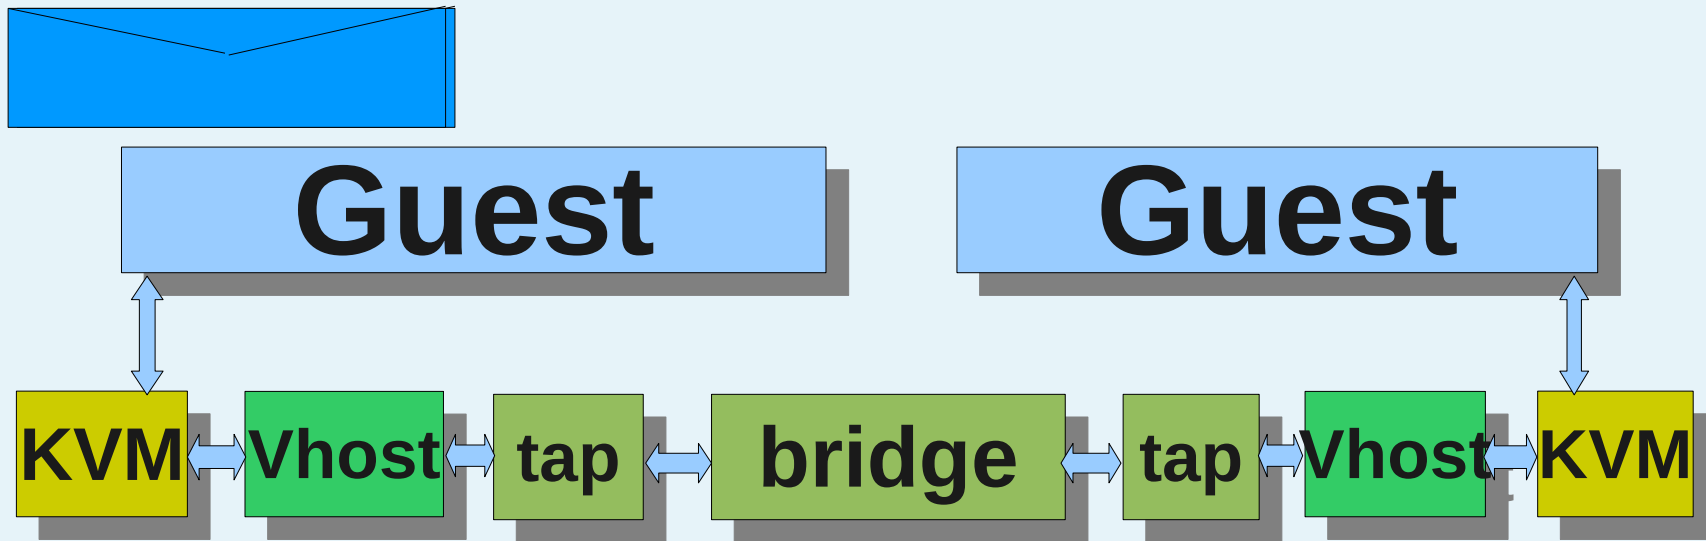
Publish guest index



Guest 2 external



Guest 2 Guest



status

- **ioeventfd/irqfd** **V**
- **Vhost-net** **V**
- **Virtio-net tuning** **V**
- **Macvtap** **V**
- **Mergeable buffers** **V**

TODO

- **Zero copy** +
- **Multiqueue** +
- **Virtio ring** +
- **Threading** ?
- **Level interrupts** ?