



**LINUX
PLUMBERS
CONFERENCE**

August 24-28, 2020



XDP guest offloading

Denis Kirjanov





LINUX
PLUMBERS
CONFERENCE

August 24-28, 2020



Introduction

- XDP support is added to many paravirtualized drivers: virtio-net, xen-netfront, hv_netvsc



**LINUX
PLUMBERS
CONFERENCE**

August 24-28, 2020



Introduction

- Some on it's way: vmxnet3



LINUX
PLUMBERS
CONFERENCE

August 24-28, 2020



Introduction

- XDP is super fast on real hardware since we're processing packets before allocating an skb



LINUX
PLUMBERS
CONFERENCE

August 24-28, 2020



XDP guest offloading

- We don't want to process each packet twice (host kernel → guest kernel), we want to avoid data copying
- And it turns out that we can get a better performance while offloading a guest XDP program to a host



LINUX
PLUMBERS
CONFERENCE

August 24-28, 2020



XDP guest offloading

- We don't want to process each packet twice (host kernel → guest kernel), we want to avoid data copying
- And it turns out that we can get a better performance while offloading a guest XDP program to a host



LINUX
PLUMBERS
CONFERENCE

August 24-28, 2020



XDP guest offloading

- Implementation challenges: eBPF maps offloading, per-cpu maps, a live migration case
- Effective eBPF maps access from a guest
- XDP_REDIRECT



LINUX
PLUMBERS
CONFERENCE

August 24-28, 2020



XDP guest offloading

- Related work:
- <https://netdevconf.info/0x13/session.html?xdp-offload-with-virtio-net>



**LINUX
PLUMBERS
CONFERENCE**

August 24-28, 2020



Thank you