

Linux netdev update

Tom Herbert <tom@herbertland.com>

Netdev conference

- October 5-7th, 2016 in Tokyo
- ~200 developers attended
- Next one will be in March 2017 in Montreal
- Over 20 papers on various topic in Linux networking
- <http://netdevconf.org/1.2/>

Some highlights...

BBR - Bottleneck Bandwidth & RTT

- New CC algorithm from Google
- Model the network by sequentially probing for RTT and bandwidth
- Nice results particularly under high loss (goodput improvements to B4)
- Questions on interaction fairness with Cubic

XDP - eXpress Data Path

- High performance networking data path
- Uses BPF for programmability
- Run in kernel, user space, device offload
- “Bare metal” CPU for forwarding and packet drop (DDOS defense)
- Several driver implementations, one NIC offload, one switch offload

Hardware offloads

- Switchdev SR-IOV
- OpenNPU
- Channel based IO to accelerate containers
- eBPF/XDP offload
- SR-IOV to implement OpenVswitch
- TLS offload

More highlights

- IPv6 Segment Routing
- Discussion on kernel bypass, impact of userspace transport stacks
- Kernel TLS implementation
- Encapsulation offloads
- ILA - Identifier Locator Addressing
- VRF - Virtual Routing and Forwarding

Thank you!