

NTP over PTP

Miroslav Lichvar

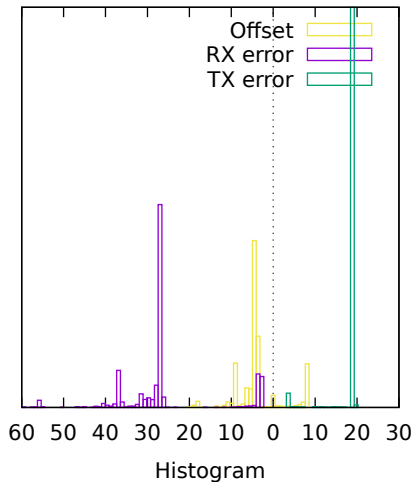

Red Hat

2021

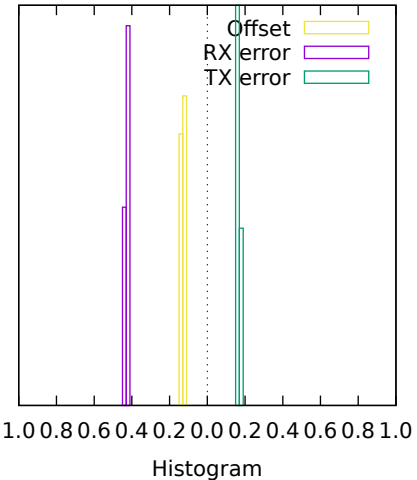

Timestamping of packets

- ▶ Accuracy, symmetry, stability
- ▶ Software
 - ▶ Application
 - ▶ Kernel
 - ▶ Driver
- ▶ Hardware
 - ▶ Physical layer (PHY)
 - ▶ Link layer (MAC)

Errors in kernel/driver timestamps



Errors in hardware timestamps



Hardware limitations

- ▶ Some HW cannot timestamp all packets
- ▶ Receive timestamping filter
 - ▶ PTP packets
 - ▶ Transport (raw ethernet, UDP)
 - ▶ Version
 - ▶ Message type
 - ▶ NTP packets (rare)
- ▶ Transmit timestamping filter (rare)

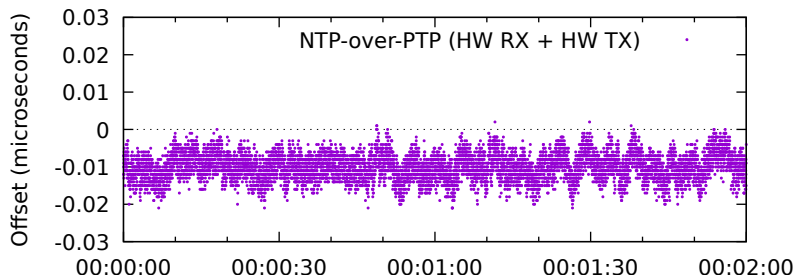

NTP over PTP

- ▶ NTP message wrapped in PTP message
- ▶ NTP gets full HW timestamping on PTP-only HW
- ▶ PTP message format
 - ▶ Version 2
 - ▶ Delay request
 - ▶ Domain 123
 - ▶ Unicast flag
 - ▶ NTP type-length-value (TLV)
- ▶ PTP correction field ignored (for now?)



PTP-only NICs tested with NTP-over-PTP

- ▶ Intel XL710 (i40e driver) - works
- ▶ Intel X540-AT2 (ixgbe driver) - works
- ▶ Intel 82576 (igb driver) - works
- ▶ Broadcom BCM5720 (tg3 driver) - works
- ▶ Broadcom BCM57810 (bnx2x driver) - doesn't work
 - ▶ Ignores unicast addresses?

BCM5720 ⇔ BCM5720



BCM5720 \Leftrightarrow switch \Leftrightarrow BCM5720



Questions?