

# **Linux Network Latency**

Herbert Xu

# Problem Definition

- Fast Gun in the West
- Start: request received
- End: response sent
- Protocol: UDP
- Processing must reach user-space

# Anatomy

- Hardware interrupt
  - 360ns
- NAPI poll
  - 1190ns
- Enter network stack
  - 1510ns
- Task wakeup

# Anatomy

- Task wakeup
  - 9720ns
- sendmsg(2)
  - 2260ns
- Exit network stack
  - 560ns
- PCI write

# Anatomy

- Idle + Sched + app 9720ns
- Driver 2100ns
- Network stack 3770ns
- Total 15590ns

# Idle Loop

- Interrupt received
- **Prepare to resume execution**
  - ACPI
  - cpuidle
  - NO HZ
- Enter scheduler
- Switch to target process

# Network stack

- SELinux netfilter hooks
- Timestamping
  - Packet socket (DHCP)
  - NTP
- IP stack

# Driver

- Interrupt mitigation
- PCI reads/writes
- Packet management



# To Do

- Reduce idle loop latency.
- Fix DHCP client.
- Optimise IP stack.
- Optimise driver.

# Questions