DNS response rate speedup using XDP

Libor Peltan • libor.peltan@nic.cz • 2020-02-08

cz.nic | CZ DOMAIN REGISTRY

Auth DNS server performance

- QPS = queries per second
- ~~Many clients~~
- Flood (DDoS) attacks:
  - mitigate = answer ‘em all
  - mostly UDP
- More hardware = $$

Auth DNS server profile

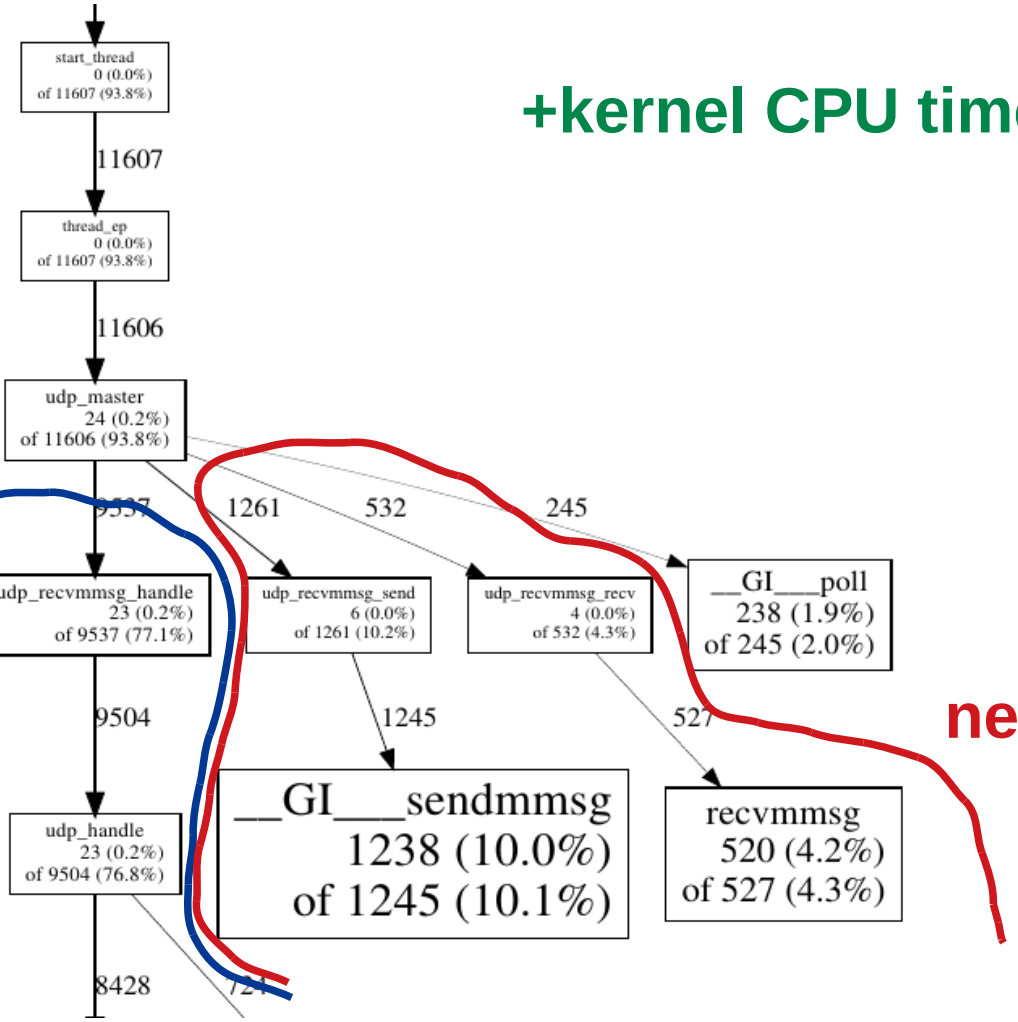

+kernel CPU time!

DNS routines

network I/O

Solution

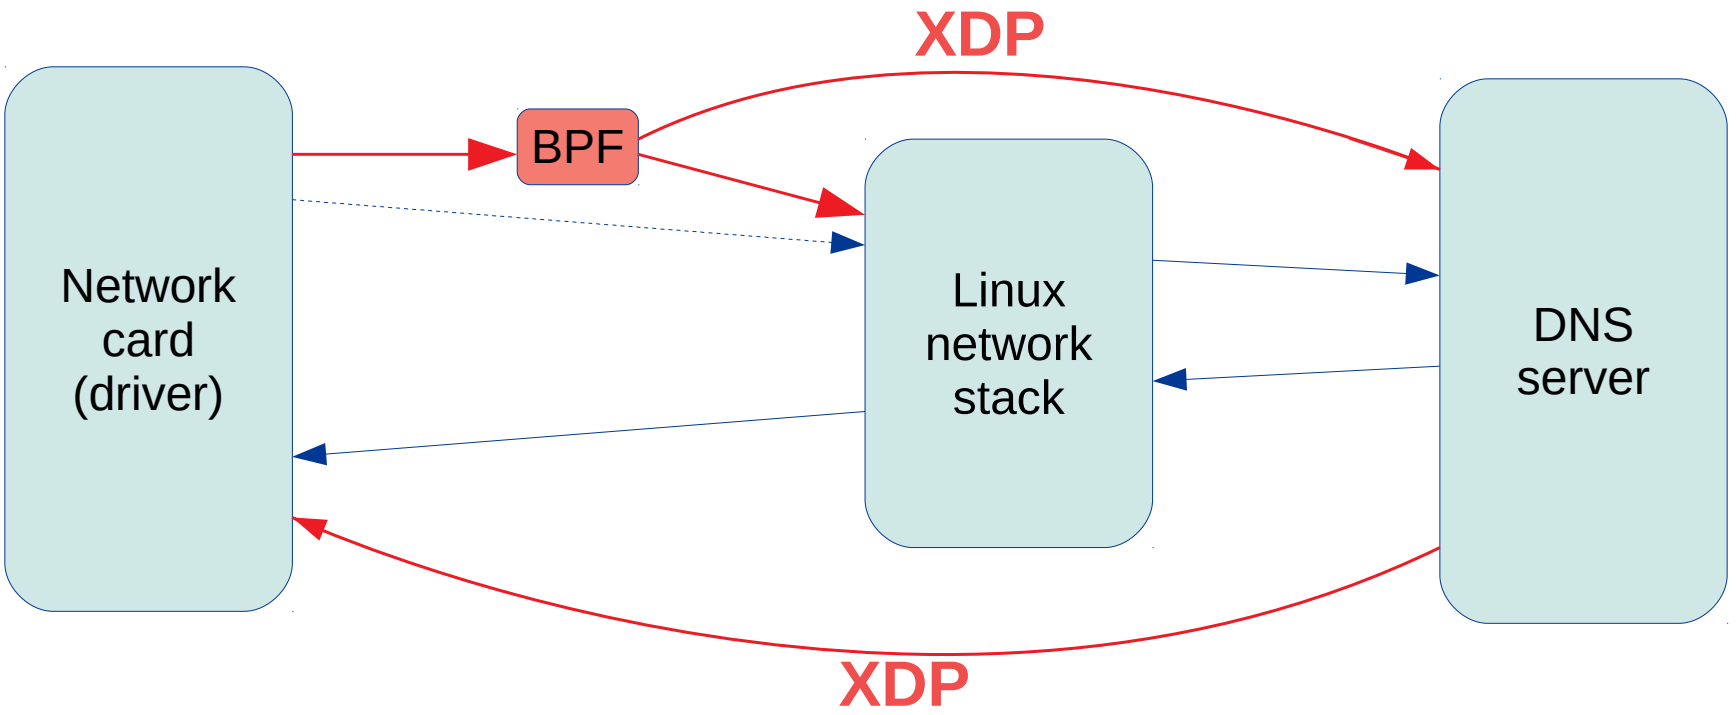
XDP
BPF
Network card (driver)
Linux network stack
DNS server
XDP

BPF (Berkeley Packet Filter)

- Originally a firewall implementation
- “BPF program” instead of rules
  - written in C
  - compiled by Clang
  - verified by kernel upon load
  - limitations (size, no loops, …)

BPF (Berkeley Packet Filter)

- BPF program decides packet fate:
  - drop
  - hand-over to XDP
    (DNS over UDP traffic)
  - pass to Linux stack
    (TCP, other port, IPv6 extensions, IPSec, etc.)

XDP (eXpress Data Path)

- Ethernet frames directly to userspace
- And back
- Zero-copy
- Need custom parsing (ethernet + IP + UDP)
- Shared umem, care about buffer allocation

Requirements

- Linux kernel 4.18+ (5.x recommended)
- XDP-compatible network card to achieve speed-up
- `CAP_SYS_ADMIN` during server startup

LibBPF

- Library for BPF and XDP stuff
- Poor presence in Linux distros
- Embedded in Knot DNS source code (for now)

Consequences

- UDP only (for now)
- Other packets processed as before (incl. TCP)
- Symmetrical routing
- **Linux firewall bypassed**

Results

- Example from some ageing hardware and very artificial setup!
  - Knot DNS + small zone:
    - Normal UDP: 1.1 Mqps
    - UDP over XDP: 3.7 Mqps (+236%)
  - Knot DNS + NSEC3 zone:
    - Normal UDP: 0.5 Mqps
    - UDP over XDP: 0.8 Mqps (+60%)

Results

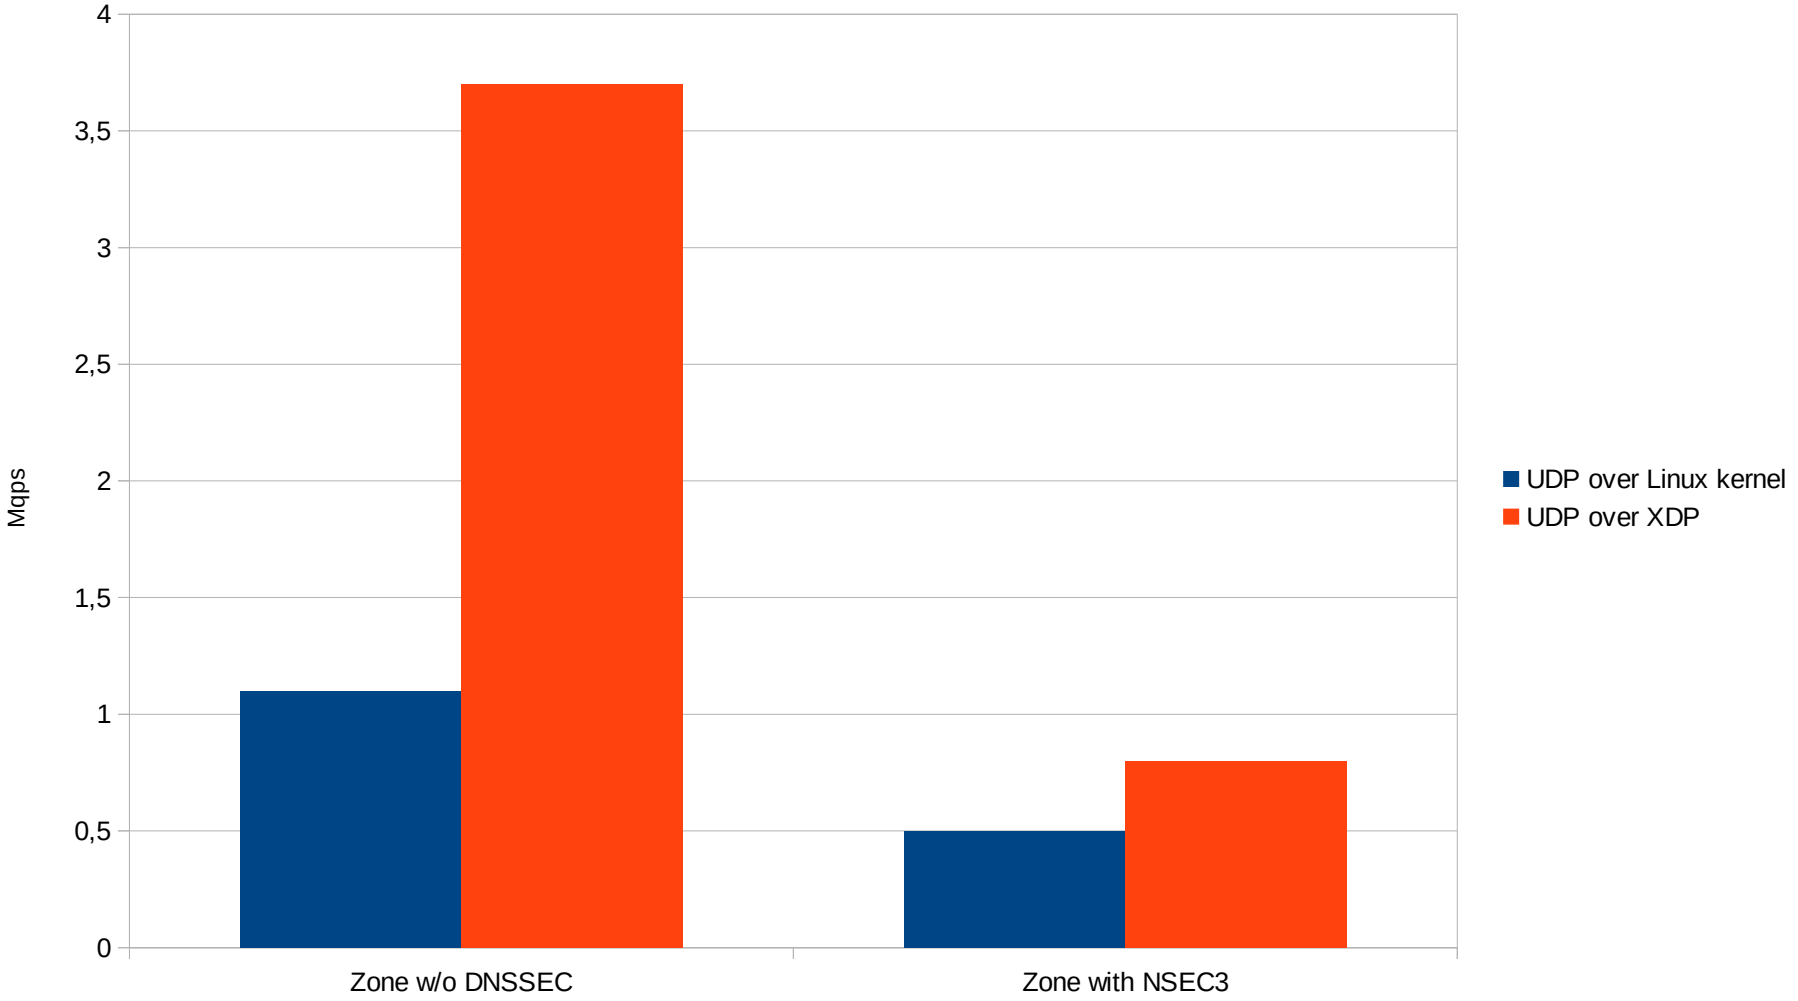
4
3,5
3
2,5
2
1,5
1
0,5
0
Mqps
Zone w/o DNSSEC
Zone with NSEC3
UDP over Linux kernel
UDP over XDP

Experience so far

- Latency also improved
- Kernel and ksoftirqd issues
- Slightly uneven CPU usage among cores

Implementation in Knot DNS

- Currently in development branch
- Experimental packages available (ask me)

- Future: also Knot Resolver

(speed up answering from cache)

Special thanks

- Vladimír Čunát, Knot Resolver