



Independent Performance Verification of XDP-accelerated Knot DNS



Brad Ellis

Solutions Engineer
brad.ellis@cengn.ca

About Me

- Brad Ellis, Solutions Engineer at CENGN
- Center of Excellence in Next Generation Networks
 - Based in Canada
 - Non-profit org supported by government and industry
 - Mission: Help Canadian startups overcome technical challenges that are inhibiting their growth

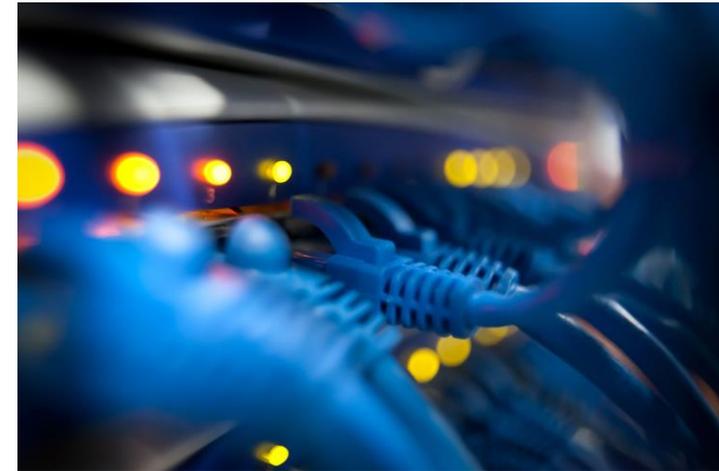
Agenda

1. Project background
2. Test Methodology
3. Baseline & First attempts
4. What is XDP?
5. XDP-accelerated Knot
6. UDP & TCP performance on XDP-accelerated Knot
7. CZ.NIC – CENGN performance comparison
8. Questions

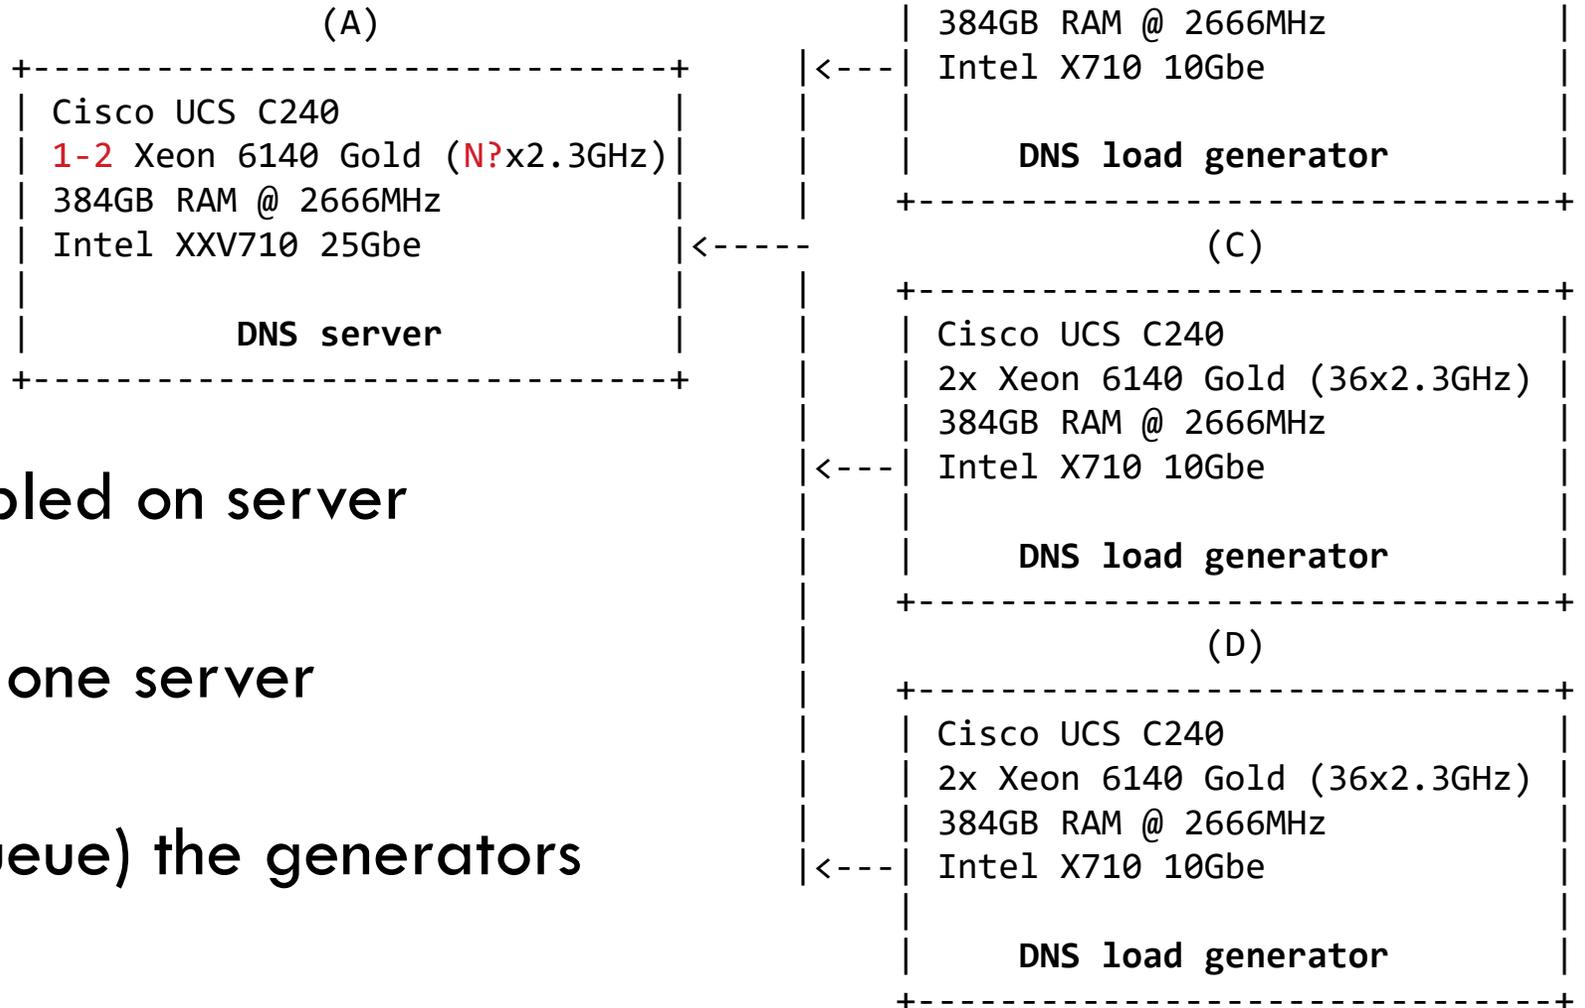
Project background

CIRA commissioned CENGN to:

- Research methods to achieve 10 Gbps UDP response rate per DNS server (or get as close as possible)
 - Document setup procedures
 - Recommend system sizing for servers
- Investigate possible DNS TCP speedup



Test Topology



- Vary # of CPU cores enabled on server
- Three load generators → one server
- Multi-thread (and multi-queue) the generators

Load Generators

- Run multiple `dnstperf/kxdpgun` instances across multiple servers
- Each load generator instance uses a dictionary derived from full zone
 - Records requested at random
 - 100 outstanding queries, after that receive response (or time-out) before sending next request
- Scripts start all load generation instances and aggregate the output of each

Baseline

- BIND9 Authoritative TLD server with real, full zone
- Baseline max UDP response of 120k QPS*
 - * DNS Queries Per Second
 - Low network utilization (168 Mbps)
 - Low CPU and Mem utilization (35% and 4GB)
- No obvious hardware bottlenecks



First attempts

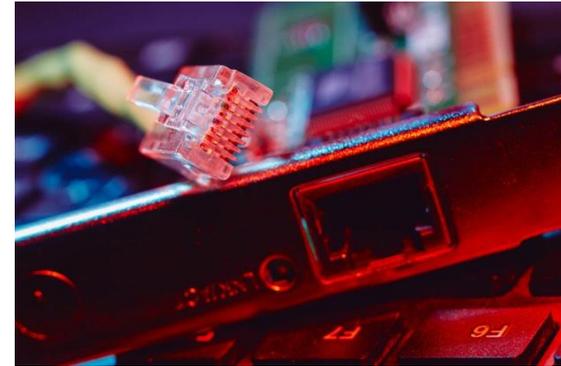
- Test several BIOS performance and power profiles
 - Land on I/O optimized profile
 - Turn hyper threading off
 - Results: ↑250k QPS (353 Mbps) ↑60% CPU
- Make cloud-native
 - Build bare-bones BIND9 pod
 - Build Kubernetes deployment
 - Results (per K8s worker): ↑515k QPS (728 Mbps) ↑90% CPU

Baseline:
120k QPS (168 Mbps) 35% CPU



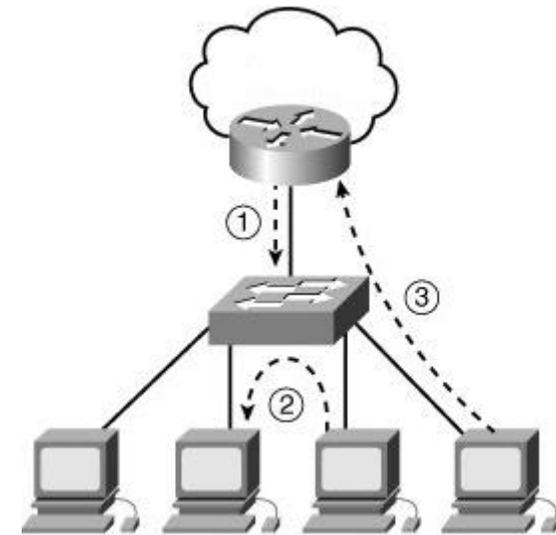
What is XDP?

- "eXpress Data Path" aims to improve network performance
- Moves many network functions lower in the 'stack'
To kernel level or even into the hardware of the NIC in some cases
- Frees up CPU and stack actions happen sooner
- Very effective in packet filtering or routing applications



XDP Limitations

- Comes with a wealth of caveats and restrictions
 - Direct access to NIC
 - Symmetrical packet paths
 - Kernel must support XDP *AND have it enabled* (CONFIG_XDP_SOCKETS=y in kernel config)
 - Application must support XDP to program it
- Very hard to add to Kubernetes
 - Hard to enforce cluster-wide symmetric packet path
 - Network plugin must do XDP in place of application in pod



XDP-accelerated Knot DNS

- CZ.NIC has been adding XDP support to Knot over the last year
- Enabling XDP in Knot's config is super easy:
xdp:
 listen: <intf>

 # optional:
 tcp: on
 tcp-max-clients: 2000000
 tcp-inbuf-max-size: 1000M
- A faster server needs a faster tester – kxdpgun
 - XDP enabled DNS load generator
 - kxdpgun multithreads to number of queues configured on NIC by default

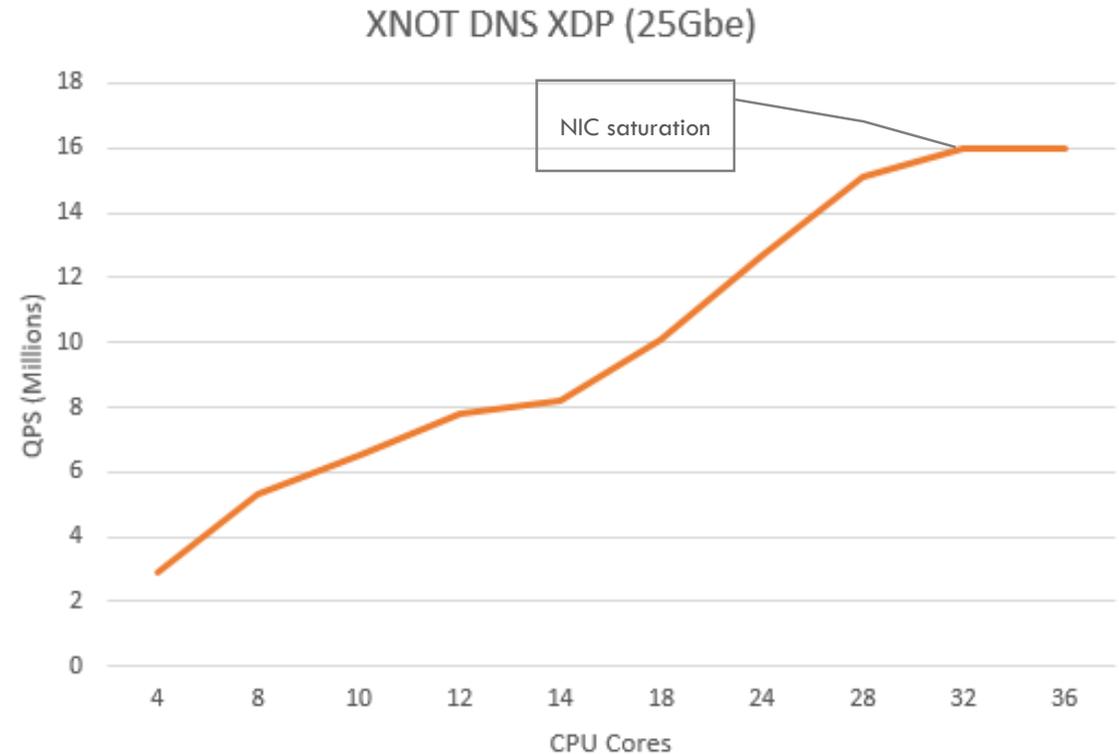


UDP Results on XDP-accelerated Knot

- ↑↑↑ 16M QPS 92% CPU
16M QPS = 25Gbps
- 25GbE NIC limited at ~32 cores
- 133.3x over baseline

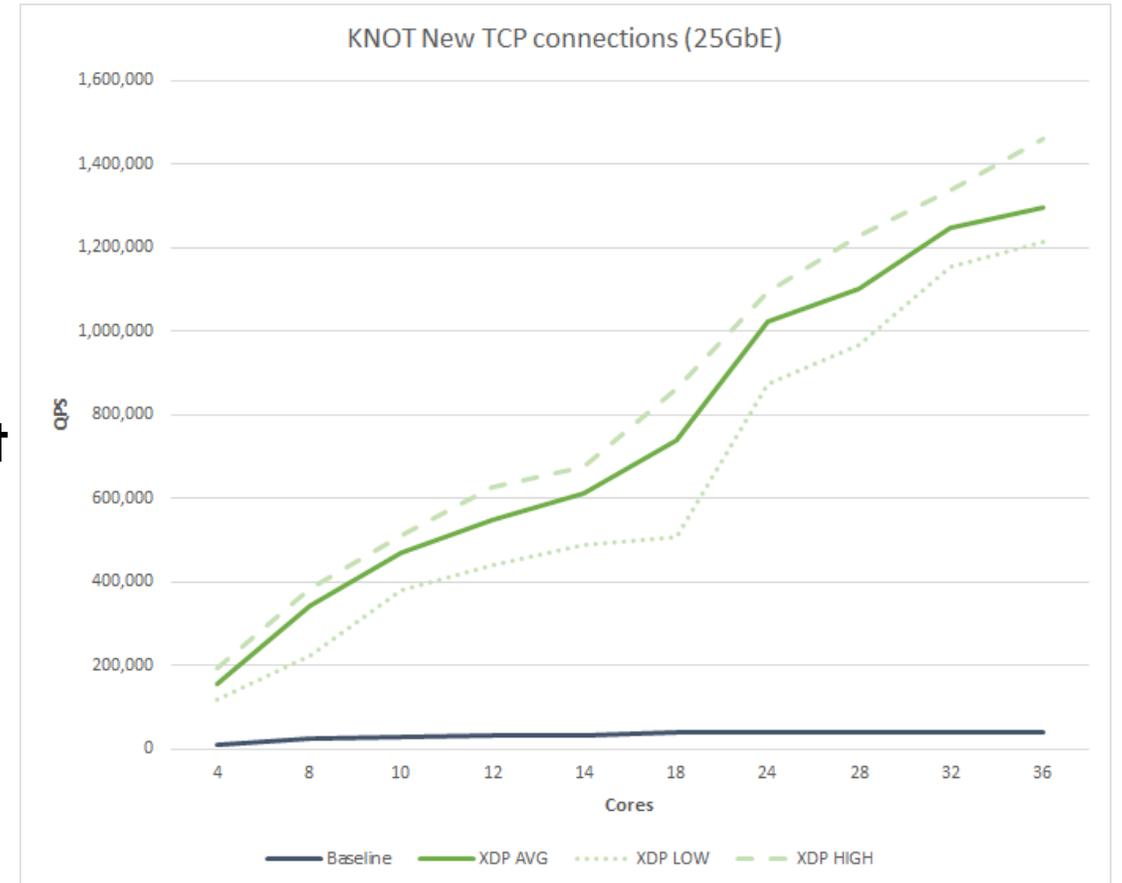
Baseline:

120k QPS (168 Mbps) 35% CPU



TCP Results on XDP-accelerated Knot

- Baseline of 40k QPS with no XDP
- Collaborated with CZ.NIC when they were developing Knot TCP XDP support
- With XDP ↑↑↑ 1.3M QPS (32.5x over baseline)



Measuring up to CZ.NIC's results

How do you compare apples and oranges?

- Different core counts, speeds and architectures
- Different NIC speeds

Best effort:

- Compare the QPS when CPU-limited and adjust for total clock cycles: **QPS/Core·GHz**
- Won't account for NUMA or architecture (in)efficiencies

	CZ.NIC	CENGN
CPU	1 x AMD	2x Intel
Total cores	64c x 2.0GHz	36c x 2.3GHz
NIC	Intel 40Gbe	Intel 25GbE
UDP QPS	22.4M	16.0M
UDP QPS/c·GHz	187k	217k
TCP QPS	2.5M	1.3M
TCP QPS/c·GHz	20.8k	17.6k

CZ.NIC performance data sourced from:
<https://www.knot-dns.cz/benchmark/>

Final Questions?

CENGN	https://www.cengn.ca/
CIRA	https://www.cira.ca/
CZ.NIC	https://www.nic.cz/
dnstperf	https://www.dns-oarc.net/tools/dnstperf
Knot DNS Server	https://www.knot-dns.cz/
XDP	https://github.com/xdp-project