

Benchmarking of authoritative DNS servers

... and DNSSEC impact assessment

Tomáš Hlaváček • tomas.hlavacek@nic.cz
DNS-OARC + NANOG65 • 3.-7. 10. 2015

Data acquisition

- Distel – derived system
- Automated measurement system
 - Build of the server software
 - Set up the server environment
 - Measure (*)
 - Create data inputs for R and generate JSON

* ... What? How?

Query generator

- Pre-created data sets – zones + pcap
- tcpreplay
 - netmap kernel API
 - The API claims to be wire-speed capable
 - DNS queries: 5-6 Mpps
 - Throttling causes bursts → We use two sources at lower speed

The variables and the results


- HW+OS (+ kernel version)
- OS tuning
- DNS server software
- Data set type → average response length
- % queries for existing vs. NXDOMAIN
- DNSSEC

- Results: Sent queries - answered queries


KNOT DNS development

- Benchmark-driven development process
- Experimental features / development methods
 - Different compilers
 - Feedback directed optimizations
 - Internal data structures
 - Different library versions and system tuning
 - Different NIC and server boards


Obvious situations...




Not so obvious situations



FreeBSD in more detail



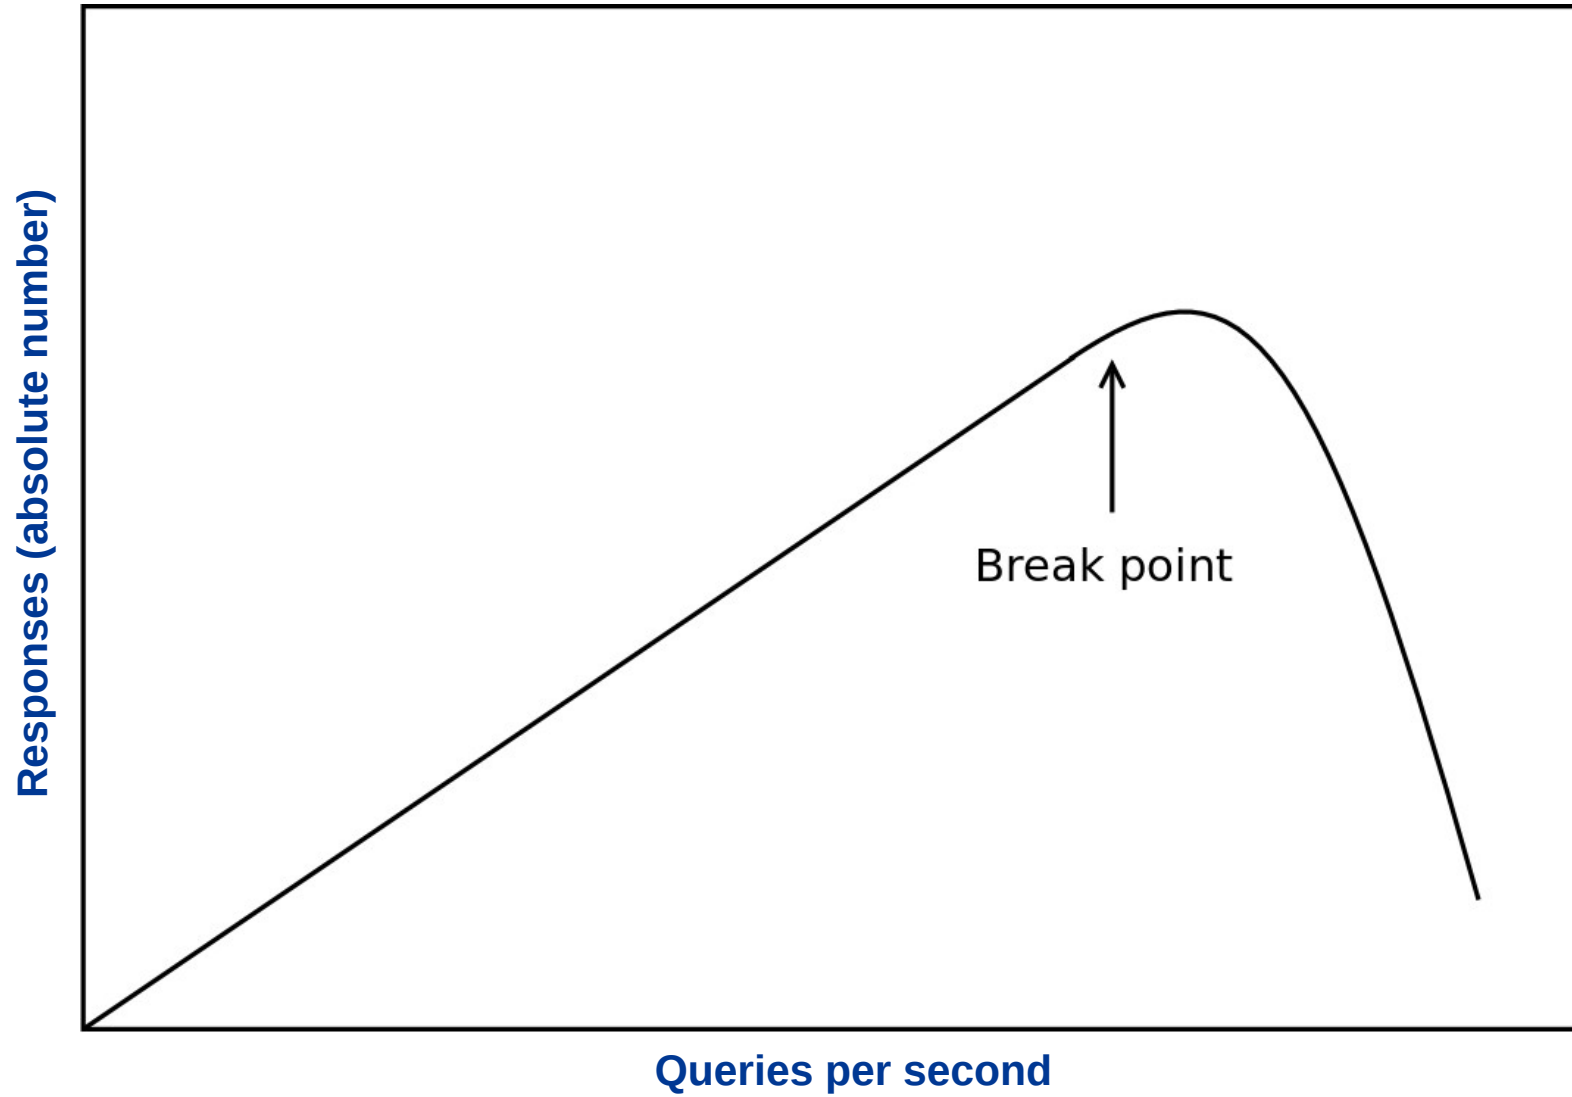
Finally statistics is needed...



Answer for developers

- Two (or more) sets of measured vectors
- Testing a hypothesis: The mean value vector is the same
 - Hotelling's test
 - We need: Same variance in both sets!
- Linear regression...

Linear model



Linear model


$$Y_i = \alpha X_i + \beta \min(0, X_i - c)^2 + e_i$$

- (Assuming Poisson distrib. of the queries)
- ANOVA to find significant factors
- Turkey HSD (honest significant difference)
- Estimate the “break point” → simple metric

DNSSEC testing


- More variables (factors)
 - Algorithm → packet length
 - NSEC/NSEC3
 - % of DO = 1 in queries
- Avg response >370 B → 3.3 Mpps max

NSEC/NSEC3 (before)



80% DO=1, 30% NXDOMAINS, algorithms 5 and 7

NSEC/NSEC3 (optimized)




80% DO=1, 30% NXDOMAINS, algorithms 5 and 7

Future work

- More automation (Ansible)
- New data sets (especially DNSSEC)
- TCP

- Comprehensive results to the web
- Discussion :-)
- dns-benchmarking@lists.dns-oarc.net ?



Thank You

Tomáš Hlaváček • tomas.hlavacek@nic.cz

Results

... in easy to read form



<https://www.knot-dns.cz/pages/benchmark.html#tab-response-rate>