



# Knot DNS

CZ.NIC, z.s.p.o.  
Ondřej Surý  
*ondrej.sury@nic.cz*  
25. 6. 2012  
ICANN 44 Tech Day

**CZ.** . . .  
**nic** :  
cz domain registry

# Design goals

- Open-source authoritative-only DNS server
  - Developed in an open way (including our mistakes)
- Usable for root, TLDs and everybody else
- Fast, feature-rich
- Portable, modular
  - Linux, \*BSD, MacOSX
  - Depend on userspace-rcu library
- More details:
  - <http://ripe63.ripe.net/presentations/145-KNOT-20111103-LS-RIPE63.pdf>

# Standards compliant

- AXFR/IXFR (both master and slave)
- All known RR Type support
  - Including TYPE#nnnn
- DNSSEC with NSEC3
- TSIG supported (from 1.0)
- Root zone support (from 1.0)
- NSID support (RFC5001) (from 1.0)
- Fast track new standards
  - DANE Protocol (TLSA RR) (from 1.0.4)

# Configuration

- Curly braces and semicolons (C-like)
  - Interfaces (IPv4 or IPv6)
  - Remotes (masters or slaves)
  - Zones
  - Keys
  - Logging (syslog or file-based)
- Runtime reconfiguration
- Precompiled zones
  - Offload the parsing from main server

# Knot DNS design

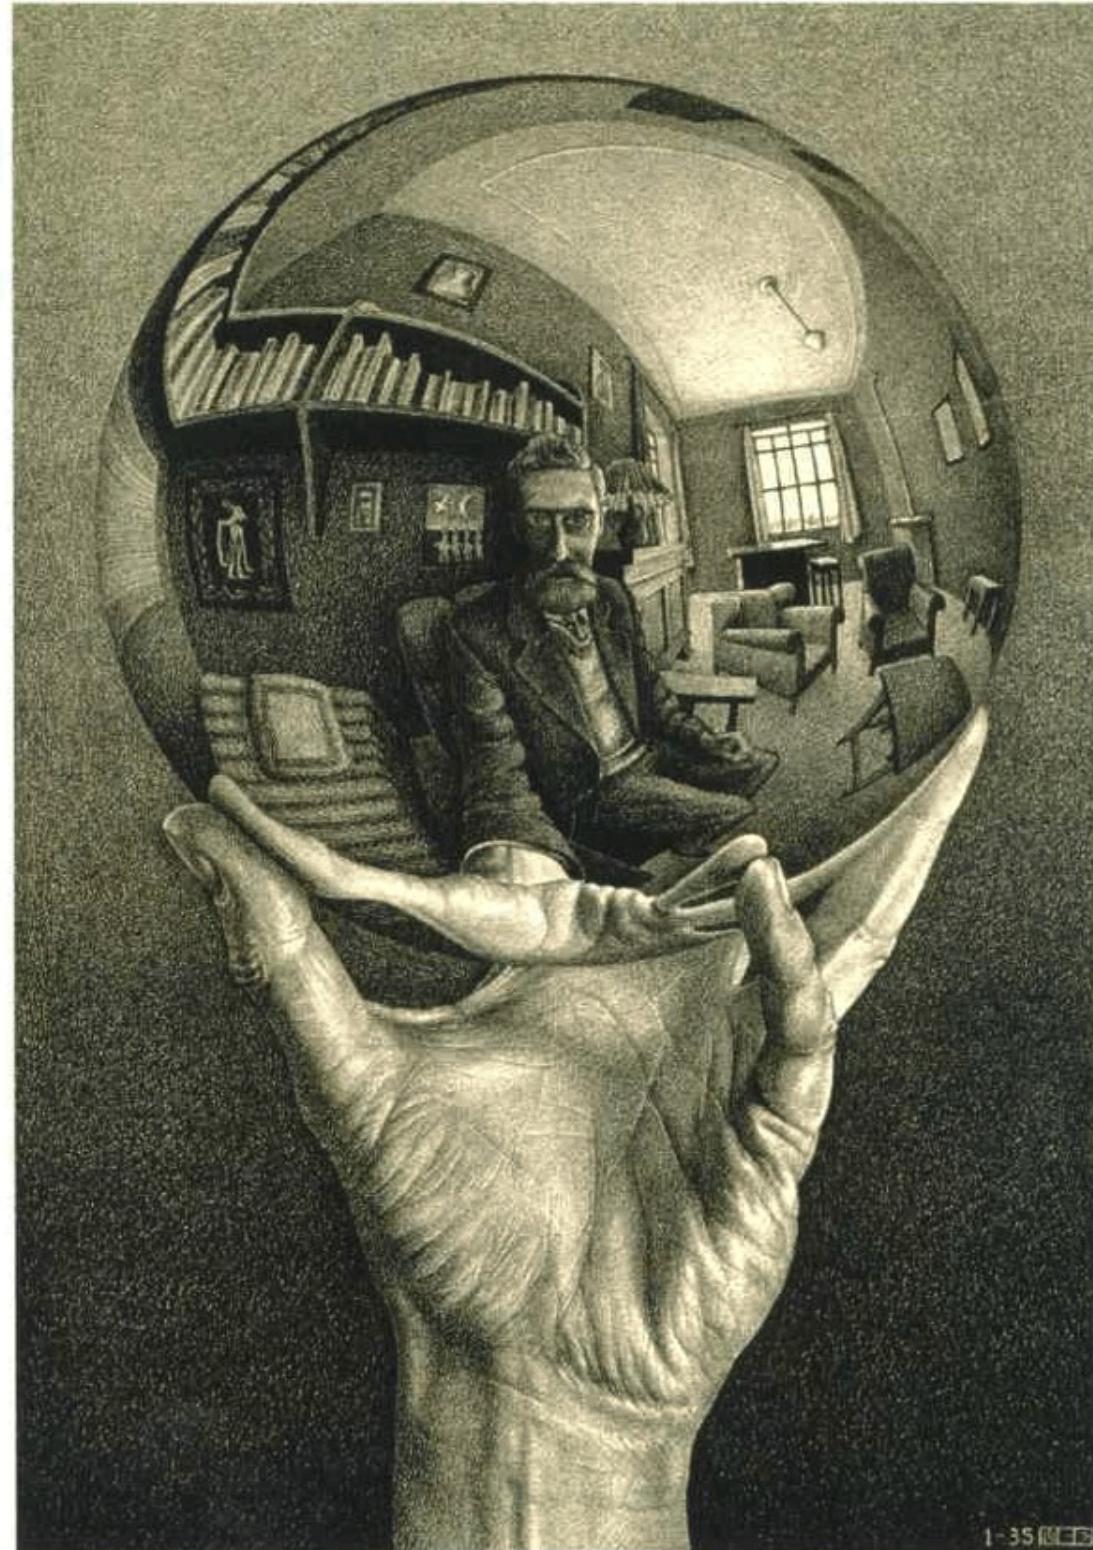
- Minimize amount of lookups for one query
  - Optimized zone structures
    - References to related data
- Minimize lookup time
  - Hash table with worst-case  $O(1)$  lookup time
    - Cuckoo hashing scheme
    - Lock-free architecture
- Non-stop operation, run-time updates
  - Read-Copy-Update (always consistent data)
  - Copy-on-Write (shallow copies)

# Roadmap

- Knot DNS 1.1 (Q3 2012)
  - Speedup of huge IXFR (40k+ records in on XFR)
  - Focus on stability and bugfixes
  - Reference Manual
    - Preliminary work (**development** branch in git)
  - Zone parsing and loading speed-up

# Future plans

- Dynamic updates
- NetConf/DNSCCM support
- Massive DNS hosting support (10-100k+ zones)
- Enhance CLI
- Your wishes?
- Talk to us :)



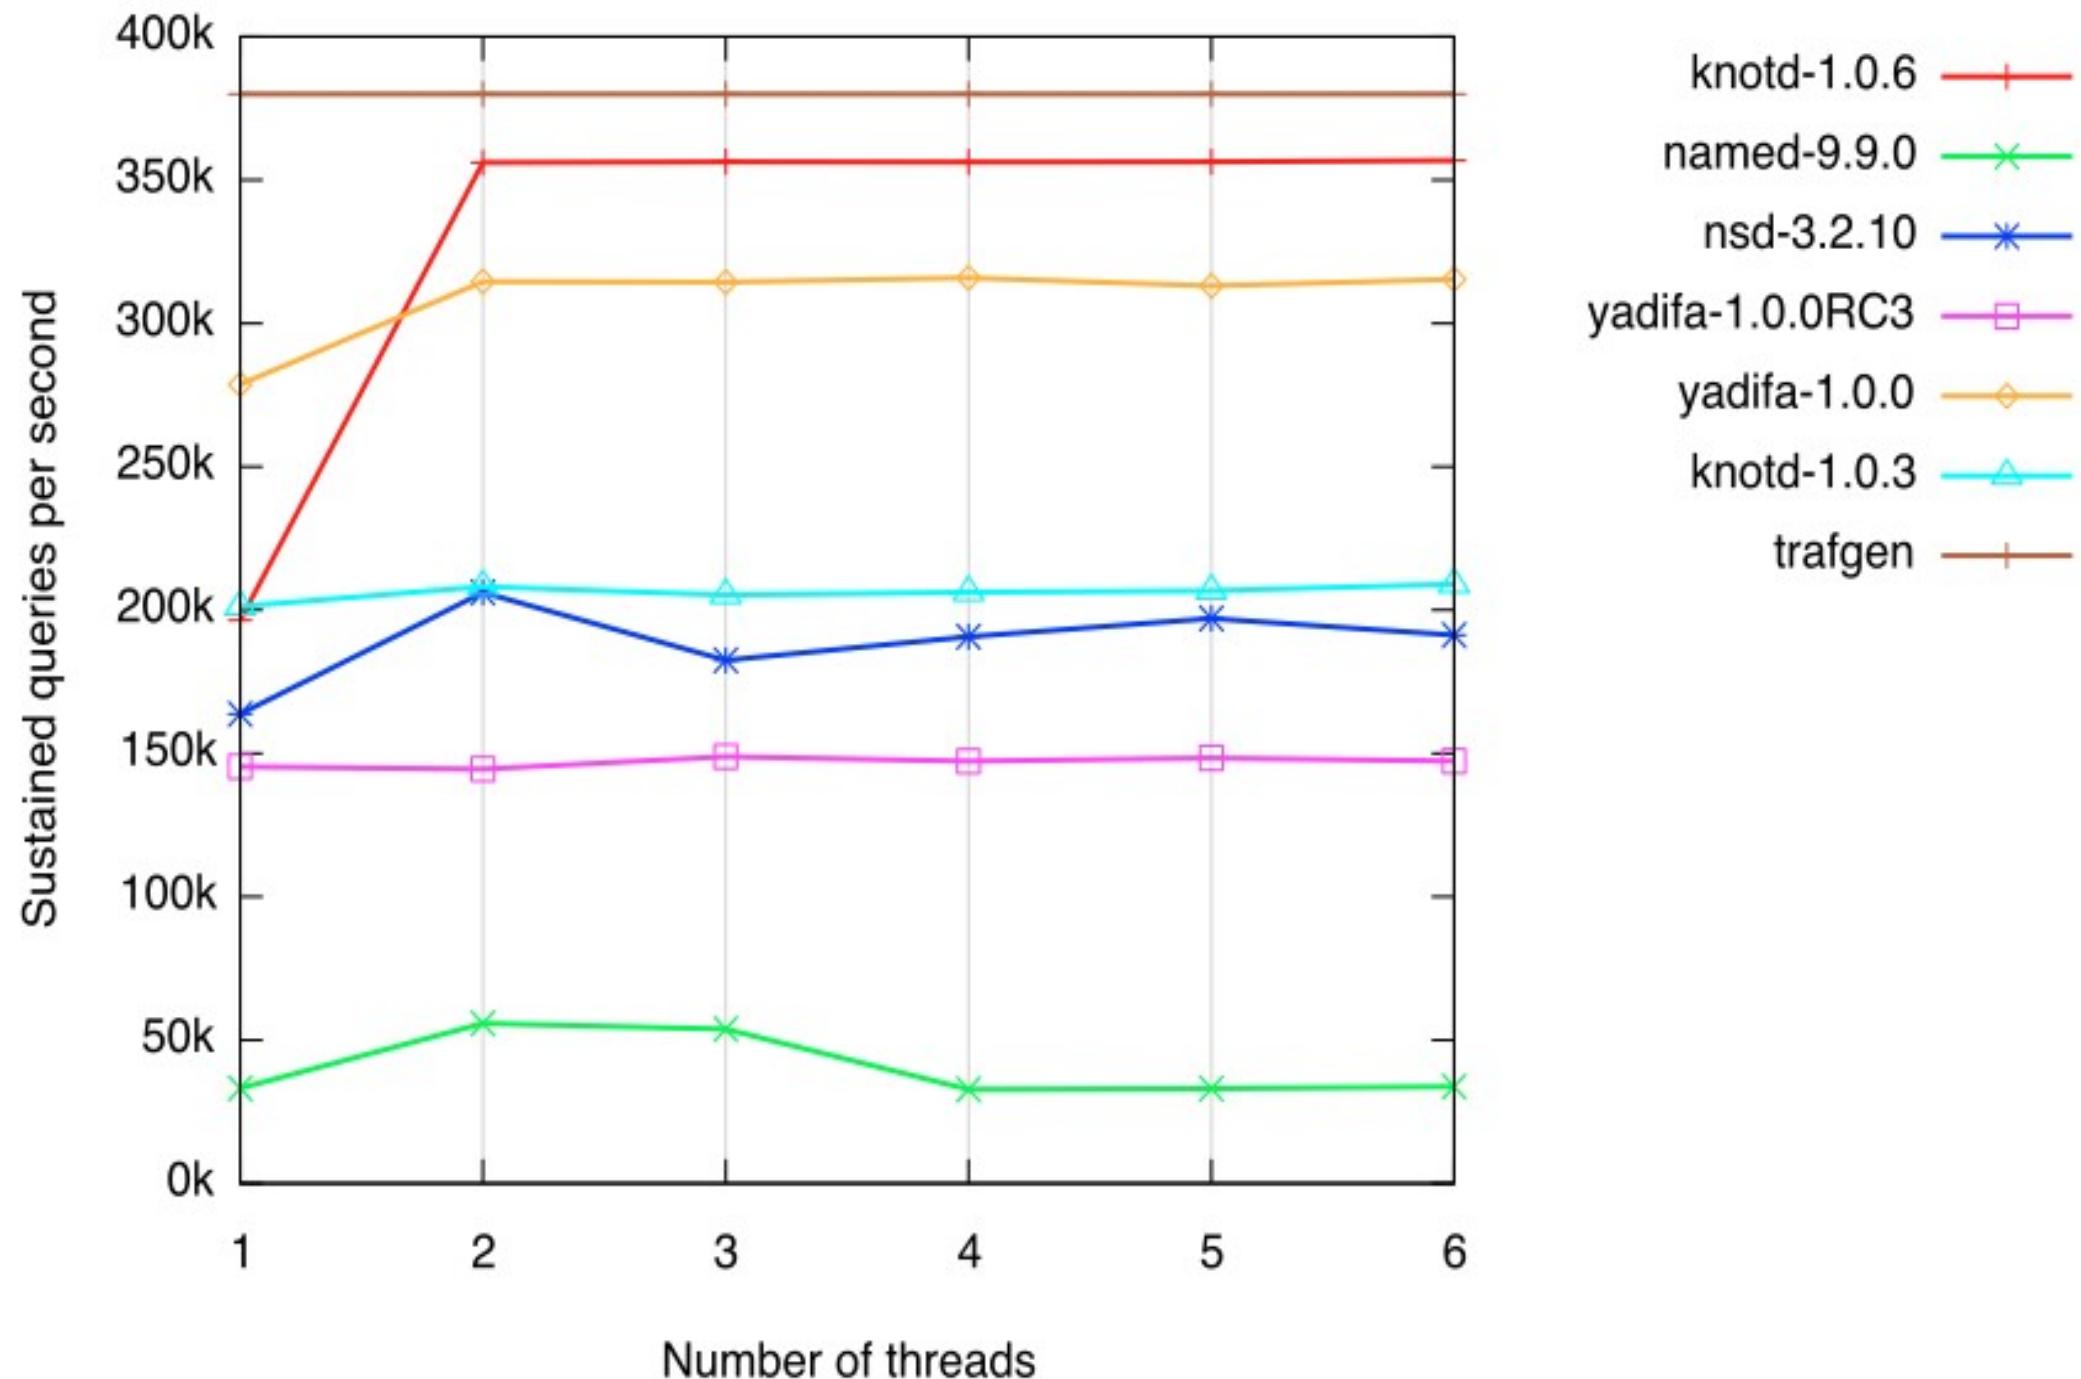
# Testing framework

- Bind 9.9.0, Knot DNS 1.0.6, NSD 3.2.10 and Yadifa 1.0.0RC2, Trafgen (<http://goo.gl/ifpKI>)
- Test zone:
  - <http://public.nic.cz/files/knot-dns/benchmark-zone.tar.gz>
  - 2 mio of random mix of unsigned records (138MB)
- Test queries
  - 50% in zone records, 50% out of the zone
  - 1 mio queries (18MB) of various type
- Commodity servers (4 Cores, 2GB)
  - Broadcom network interface

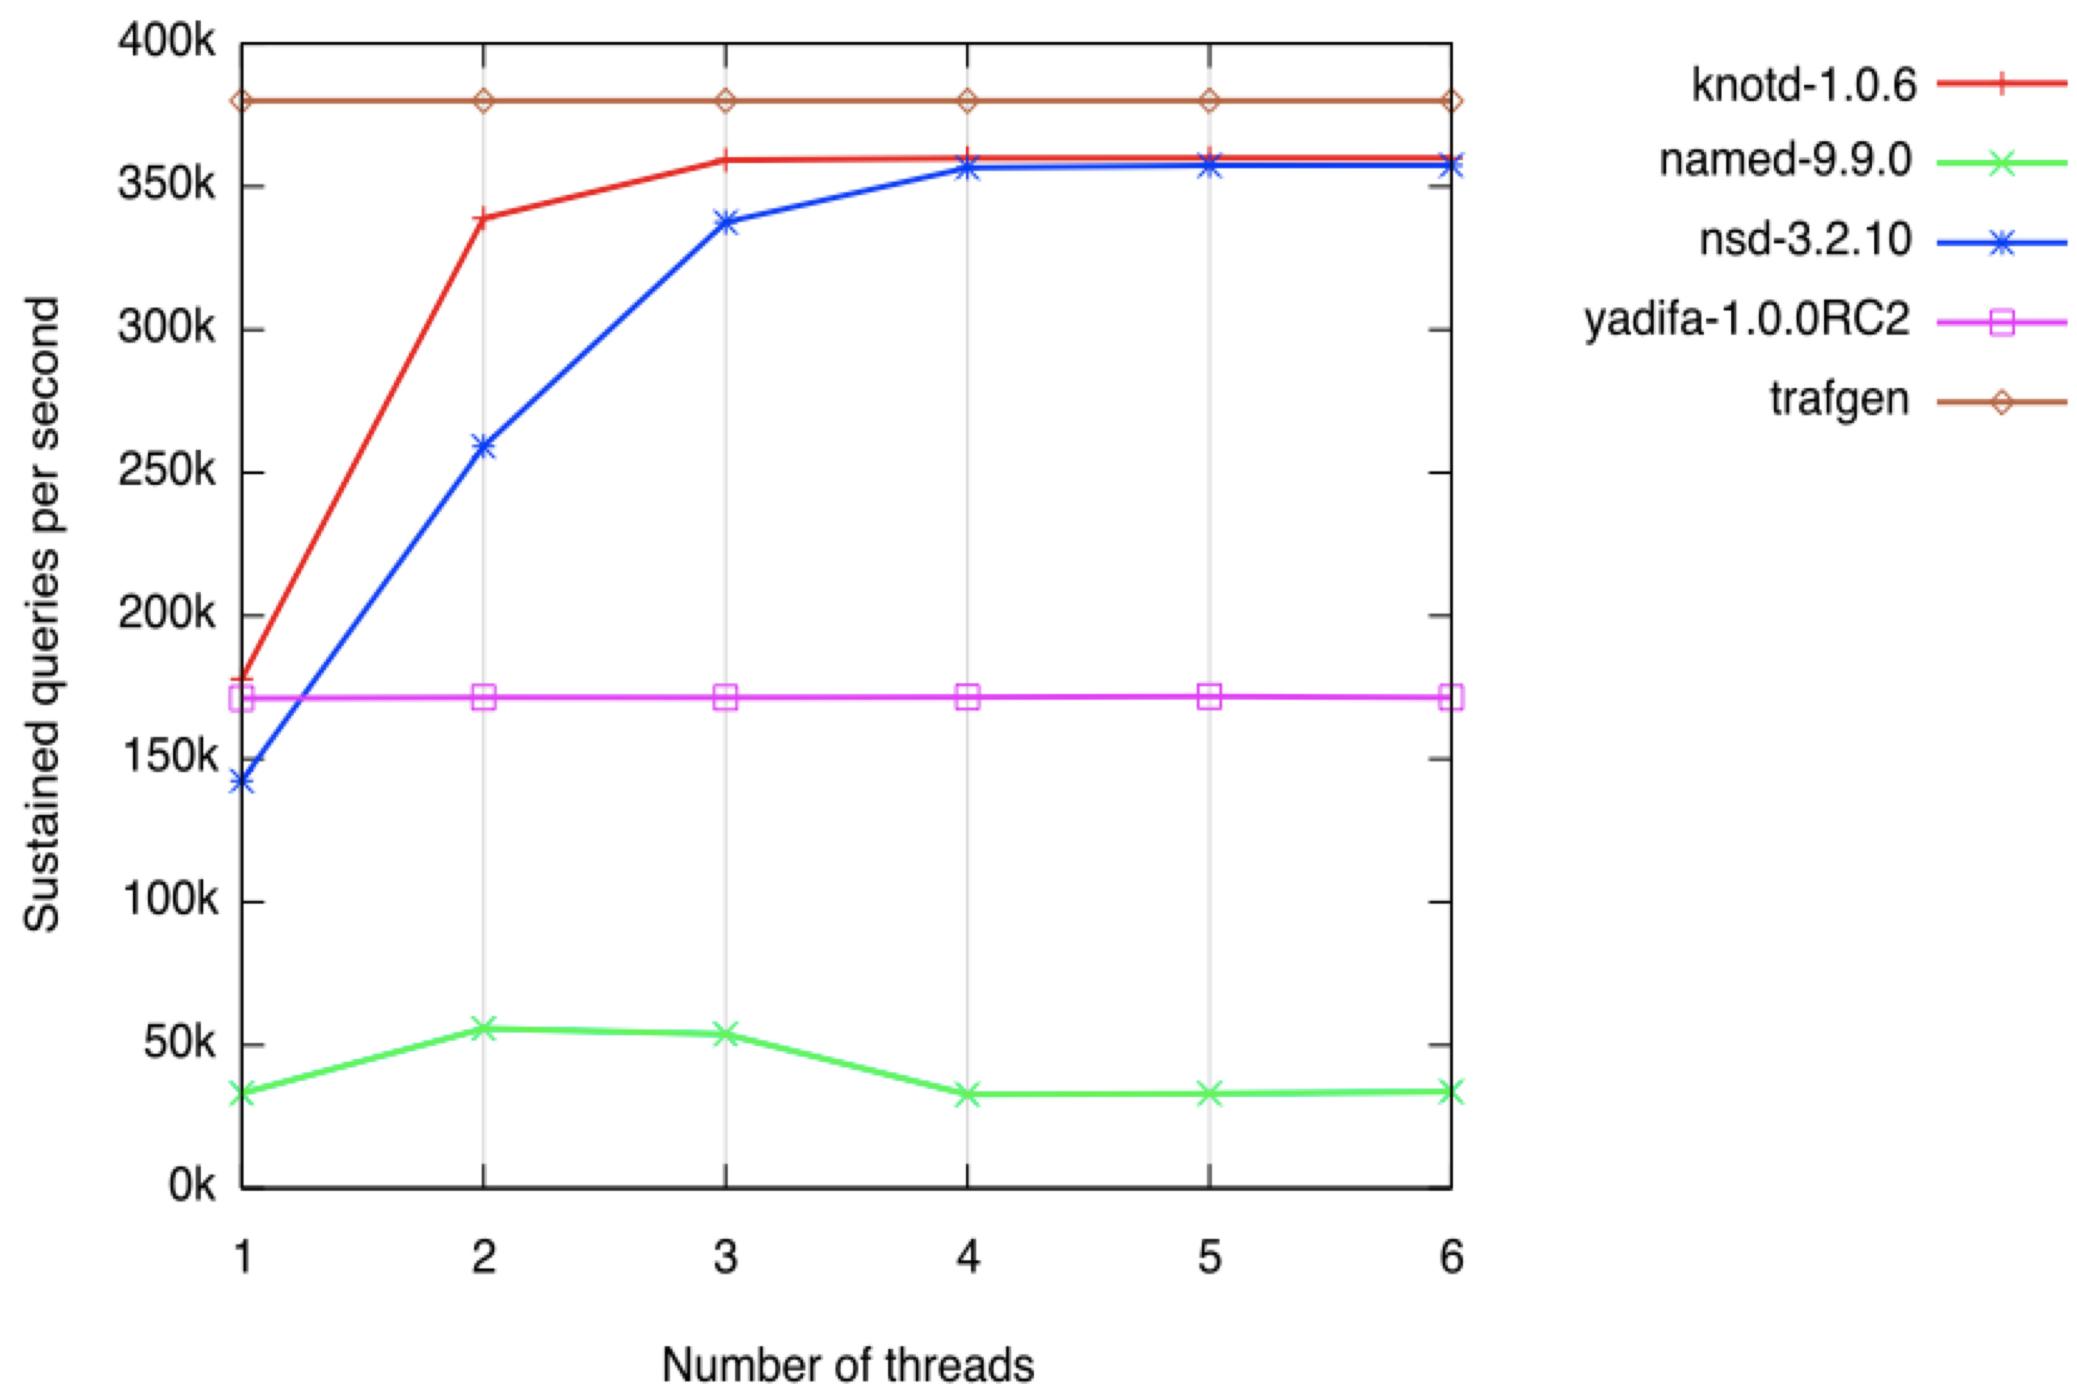
# Performance testing 1

- dnsperf based, one client per core, one server
  - Sliding window
- More iterations to stabilize the results
- Independent variable: **threads/processes**
  - Note: Yadifa has default number of threads
- Dependent variable: **queries per second**
- Two runs:
  - Linux 3.x
  - FreeBSD

### dnsperf benchmark (Linux-3.0.0 amd64)



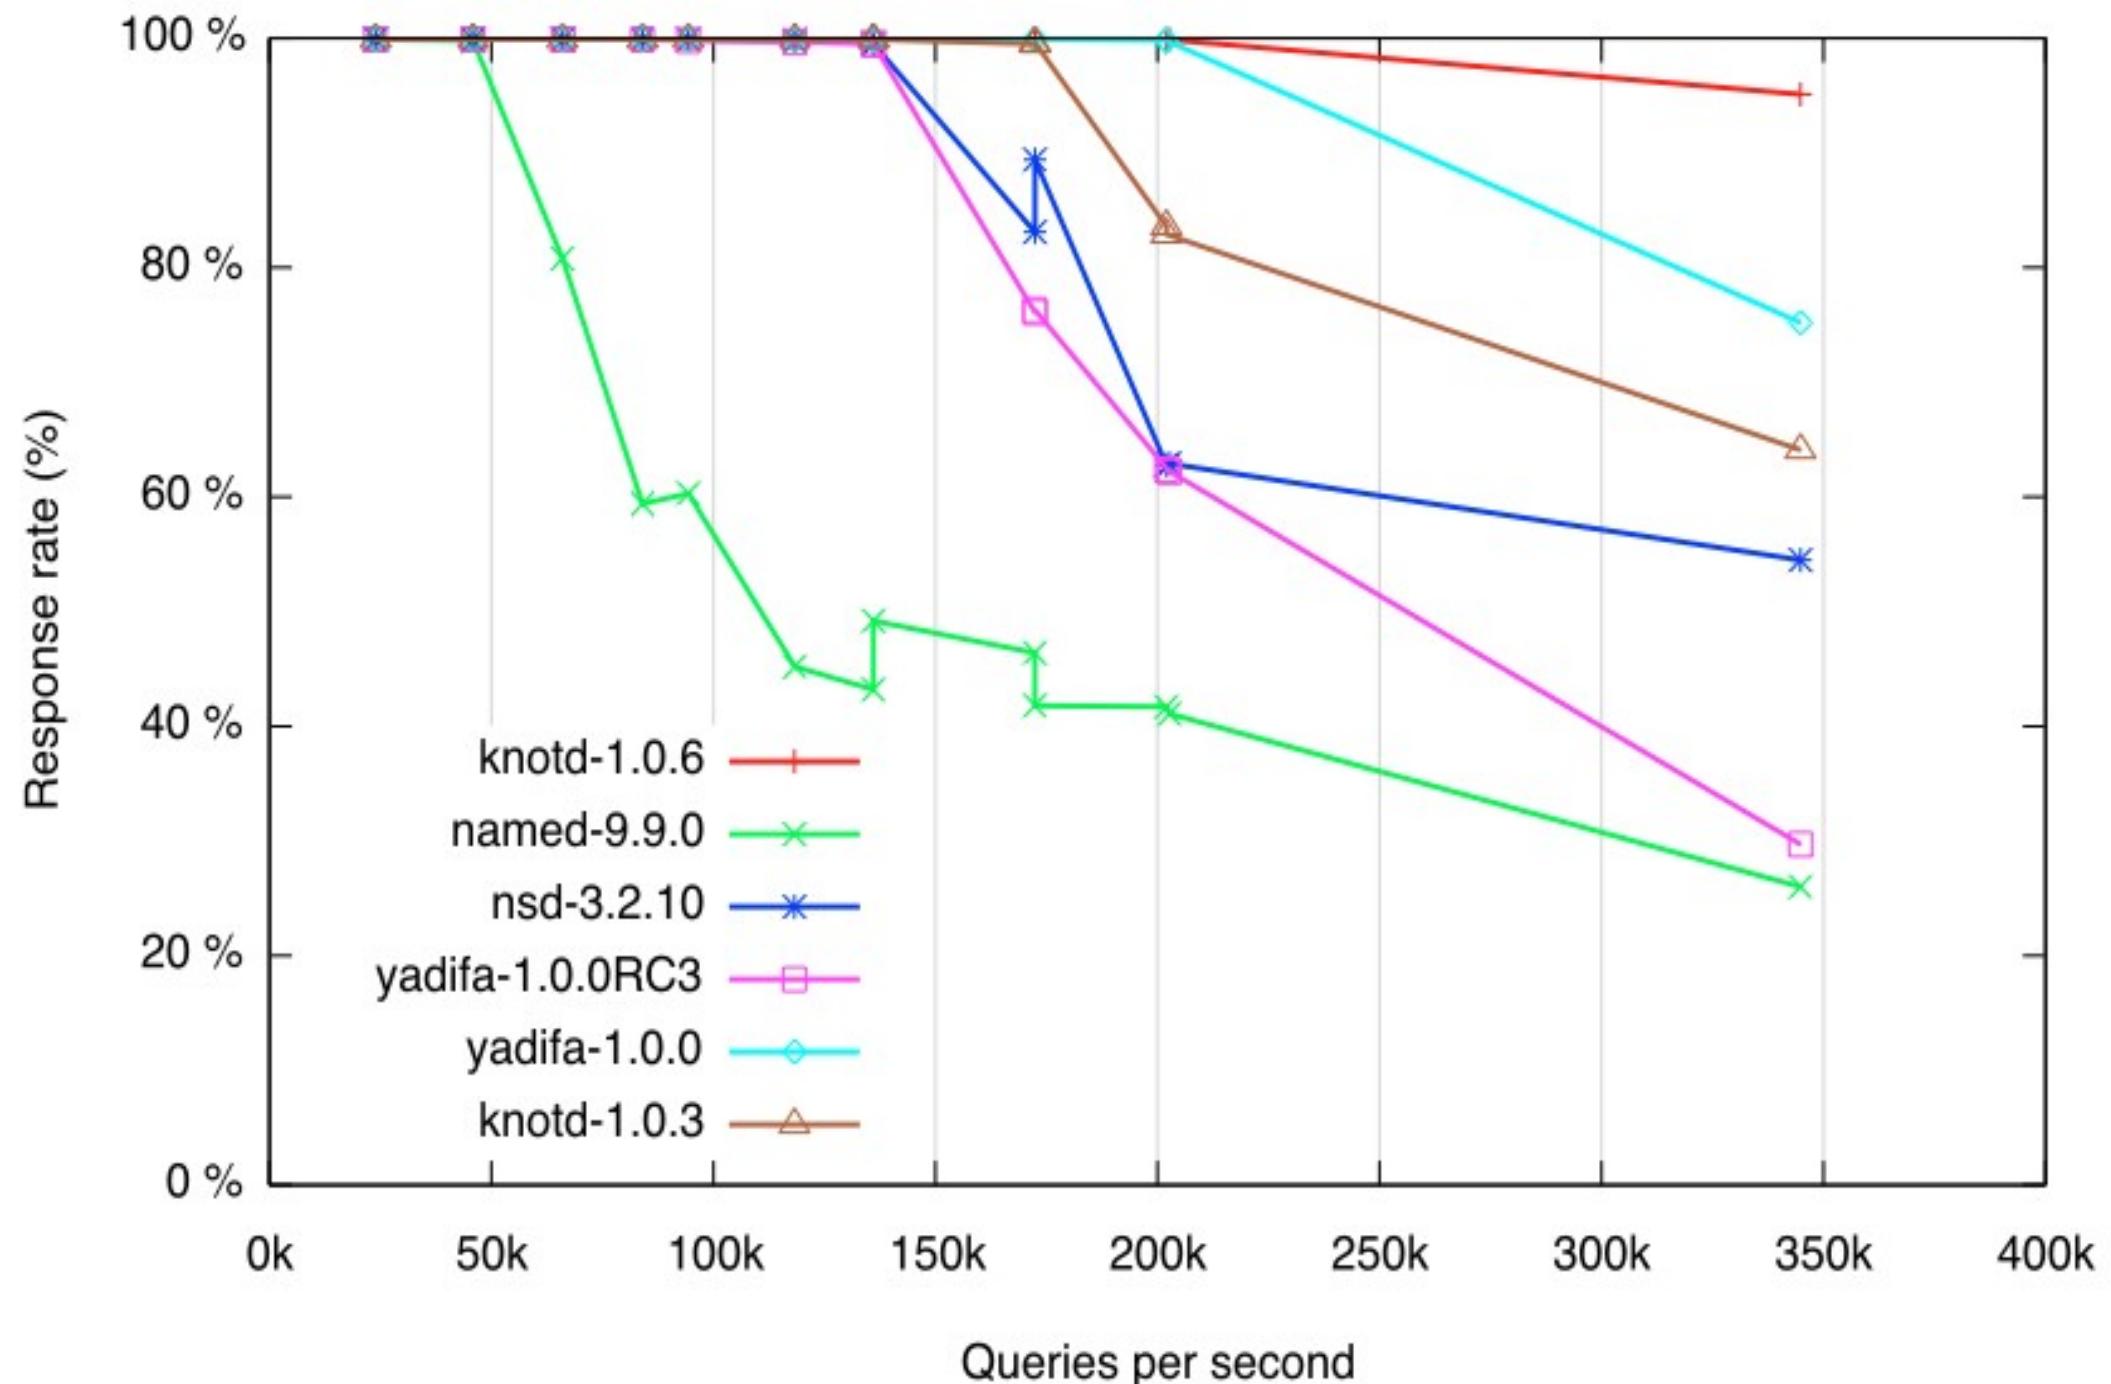
### dnsperf benchmark (FreeBSD-8.2 amd64)



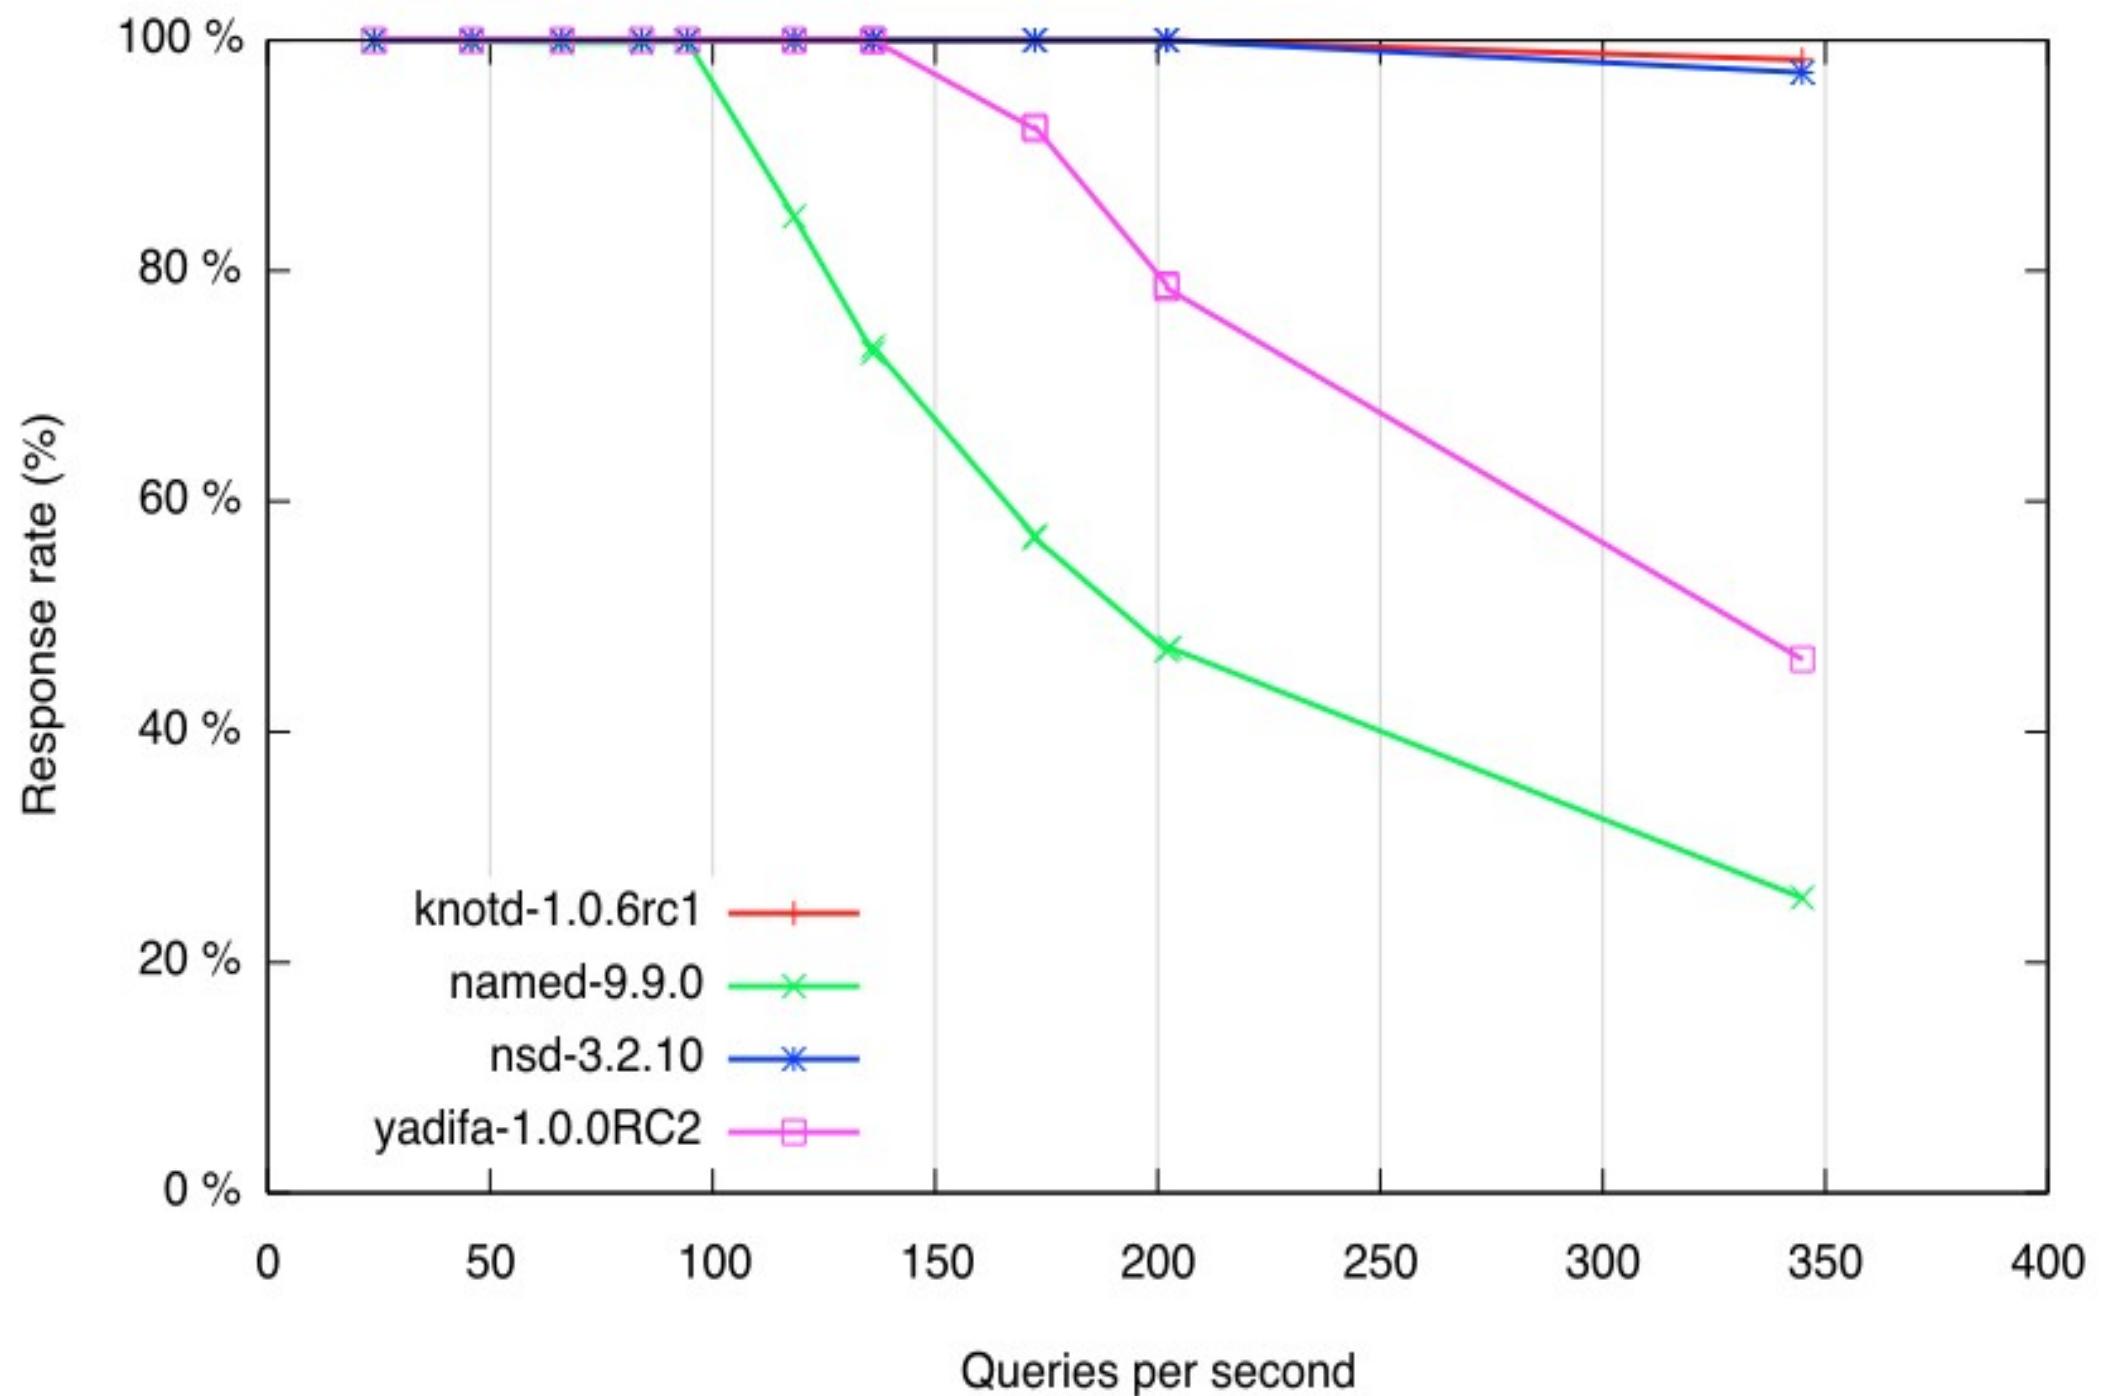
# Performance testing 2

- pcap/tcpreplay based
  - <http://www.yadifa.eu/benchmark>
- Independent variable: **queries per second**
  - Last value: --top-speed
- Dependent variable: **percentage of lost queries**
- Two runs:
  - Linux
  - FreeBSD

### Response rate (Linux-3.0.0 amd64)



### Response rate (FreeBSD-8.2 amd64)



# Pre-packaged Knot DNS

- Linux
  - Debian
    - <http://packages.debian.org/knot> (wheezy,sid)
    - deb <http://deb.knot-dns.cz/debian/> squeeze main (squeeze)
  - Ubuntu
    - <http://packages.ubuntu.com/knot> (quantal)
    - ppa:cz.nic-labs/knot-dns (lucid,oneiric,natty,precise)
  - Fedora (official packages will be available shortly)
    - <http://rpm.knot-dns.cz/redhat/>
- FreeBSD
  - <http://www.freebsd.org/cgi/cvsweb.cgi/ports/dns/knot/>

# Resources

- Home page: <http://www.knot-dns.cz/>
- Google+ page with news: <http://goo.gl/f7lWF>
- Issue tracking and source code
  - Contributions welcome!
  - <http://git.nic.cz/redmine/>
  - <git://git.nic.cz/knot-dns>
- Mailing list  
[knot-dns-users@lists.nic.cz](mailto:knot-dns-users@lists.nic.cz)