Knot DNS

CZ.NIC, z.s.p.o.
Marek Vavruša
marek.vavrusa@nic.cz
26. 9. 2012

cz nic
cz domain registry

Introduction

- Open-source authoritative-only DNS server
- Usable for root, TLDs and everybody else
- Linux, *BSD, Mac OS X
- AXFR, IXFR, TSIG, DNSSEC with NSEC3
  - and much more


THE FLASH

New features since version 1.0

- **User manual**
- Faster IXFR, faster server load
- Creating differences from zone file changes

  = Full master support
- Control utility - checkconf, checkzone
- DANE protocol support (RFC 6698)

Features planned for v1.2 release

RFC2136 Dynamic DNS update

- Secured with TSIG
- Including update forwarding
- Does not interrupt server operation
- No update consolidation so far

New zone file parser

- Rich features
  - All major RR types
  - Custom RR types with binary data
  - $ORIGIN, $TTL, $INCLUDE
- Czech zone parsed in 3 seconds
- Completely open-source, Knot independent
  - May be used for other DNS libraries
  - See branch **zscanner** in our repository

Remote control utility

- Secured protocol, one-way hashing
- Extensible commands, may include data
  - reload
  - refresh zones
  - flush zones
  - possibly dynamically add/remove zone
  - … and more

Knot DNS on a diet

Knot DNS on a diet

- Goal is to reduce memory usage
  - Spikes during transfers
- Consolidating internal data structures
  - Get rid of pointers
- Reducing number of active data structures
- Custom allocator for zone data

Why should you give it a try?

A: Thorough testing process

- Mostly automated
  - Around 40 scenarios tested nightly
- Static code analysis tools
- Prior to each release:
  - Code freeze
  - 2 weeks in a semi-production environment

B: We talk to our users

- User reported issues have a priority
- We love your feedback, even negative
- Google+, Twitter
- Mailing list

  knot-dns-users@lists.nic.cz
- Issue tracking system
  - http://git.nic.cz/redmine/

C: Current users of Knot DNS

- L-root (experimenting)
- Czech hosting companies
  - hosting90.cz, igloonet.cz
- CZ.NIC
  - Slave server under heavy load
  - All of CZ.NIC's zones, including **.cz**
- **.ru** (slave server, testing phase)

Recap of our plan

- Four new features for Knot DNS 1.2
  - DDNS update including forwarding
  - Blazing fast zone file parser
  - Remote control utility
  - Reduced memory consumption
- Coming November 2012

www.knot-dns.cz

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

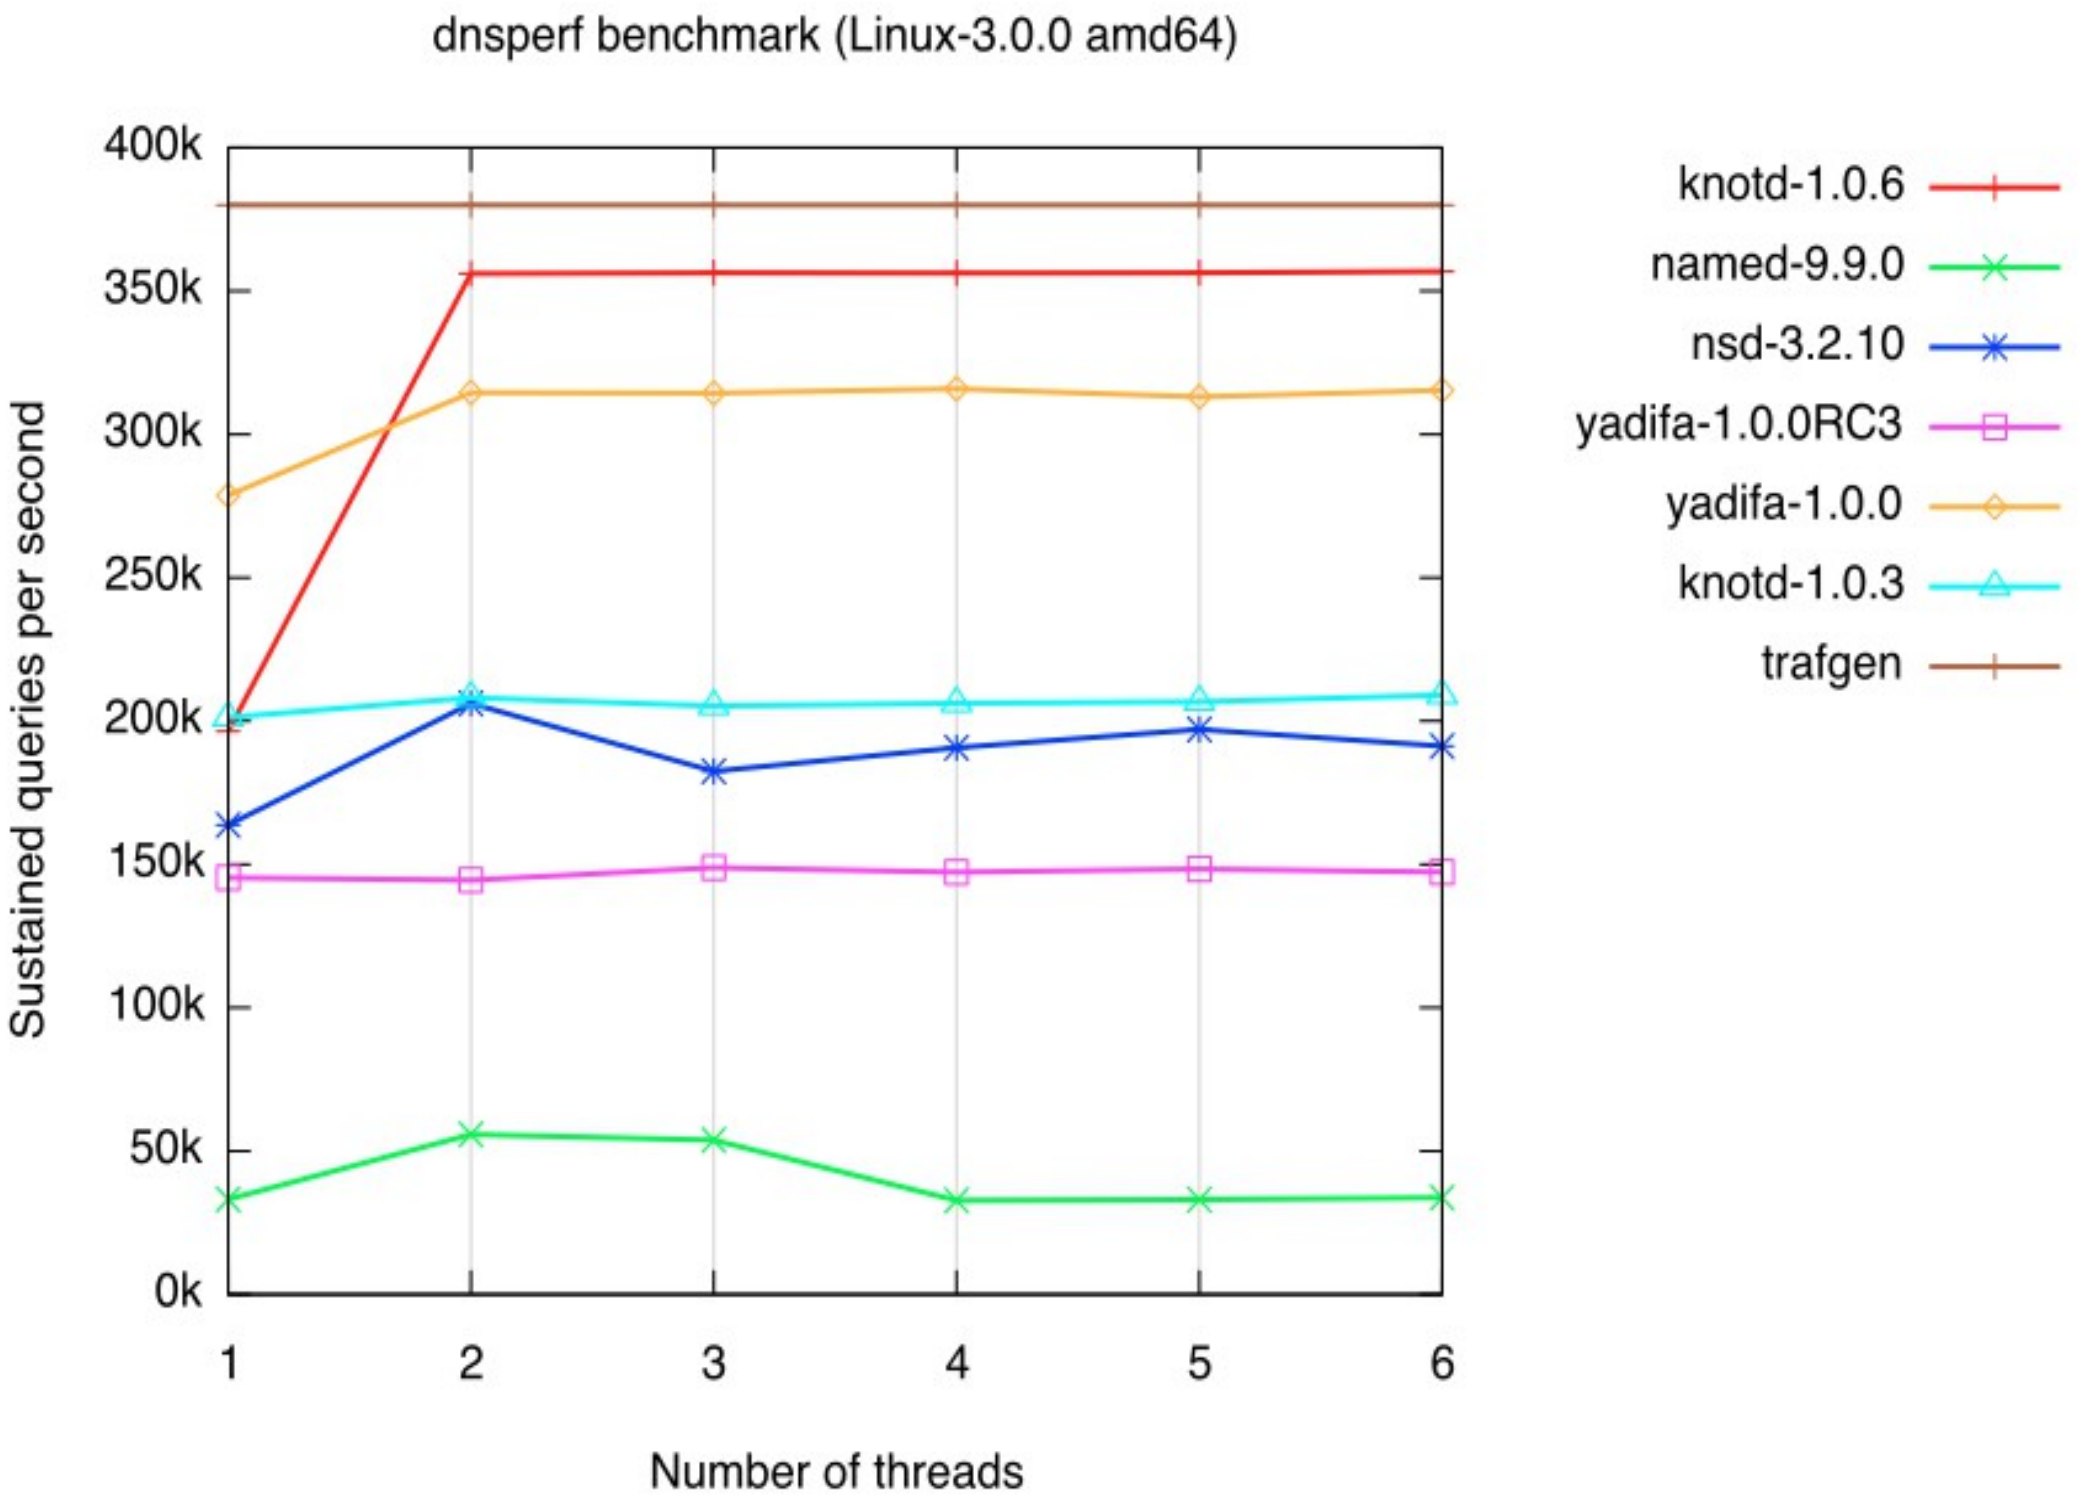
dnsperf benchmark (Linux-3.0.0 amd64)
Sustained queries per second
400k
350k
300k
250k
200k
150k
100k
50k
0k
1
2
3
4
5
6
Number of threads
knotd-1.0.6
named-9.9.0
nsd-3.2.10
yadifa-1.0.0RC3
yadifa-1.0.0
knotd-1.0.3
trafgen

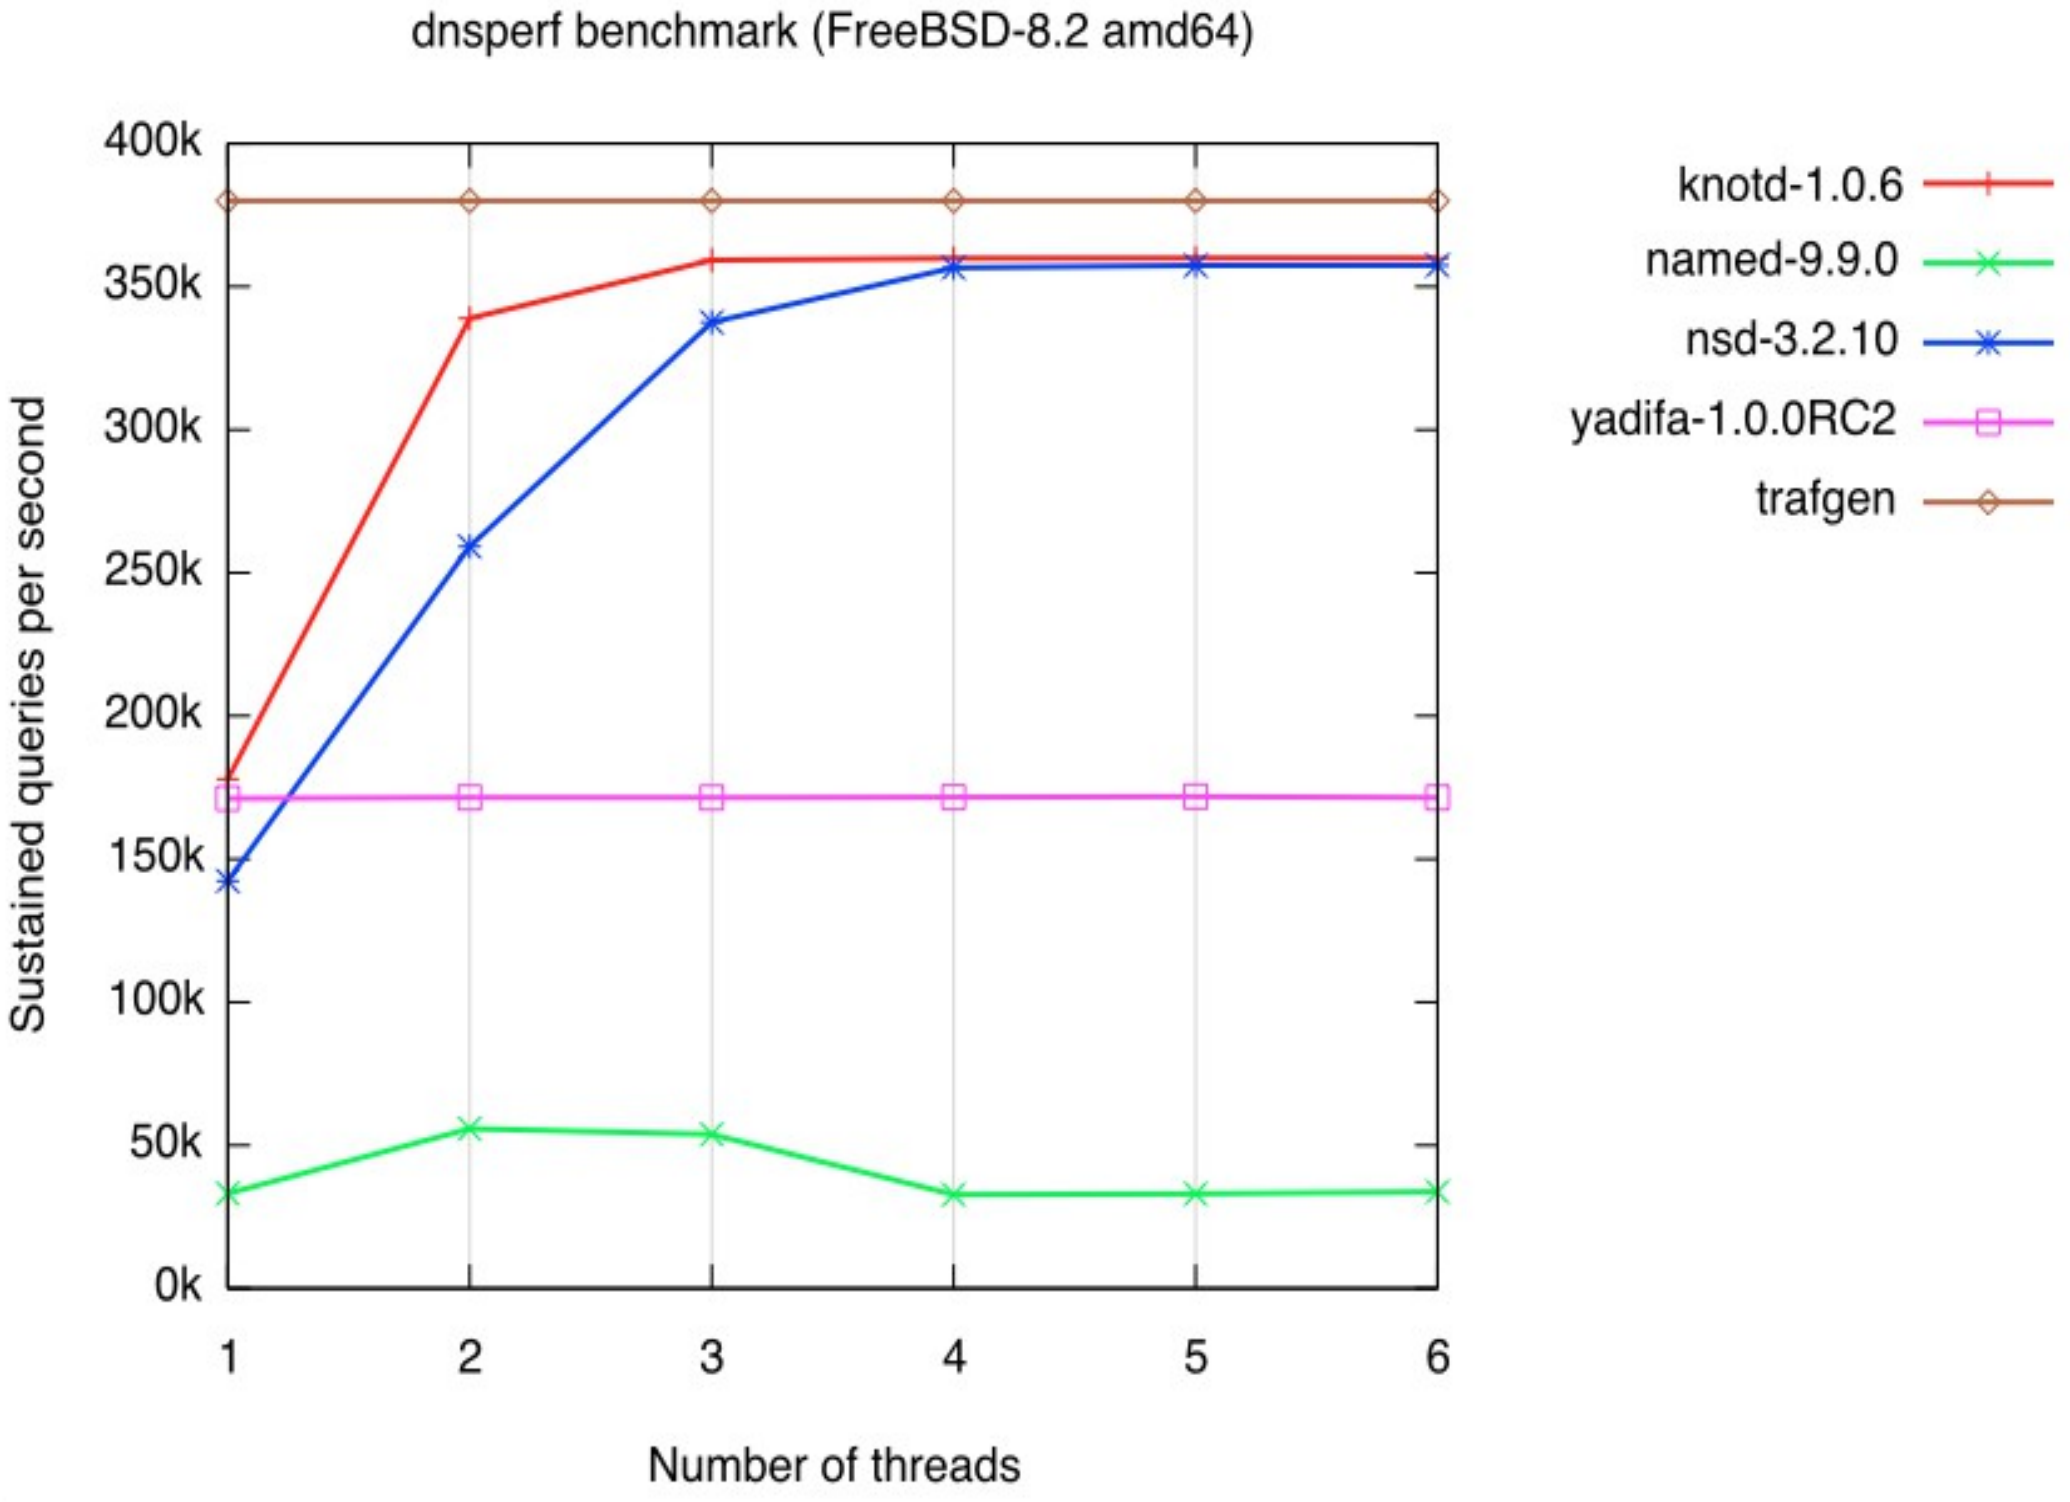
dnsperf benchmark (FreeBSD-8.2 amd64)
Sustained queries per second
400k
350k
300k
250k
200k
150k
100k
50k
0k
1
2
3
4
5
6
Number of threads
knotd-1.0.6
named-9.9.0
nsd-3.2.10
yadifa-1.0.0RC2
trafgen

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

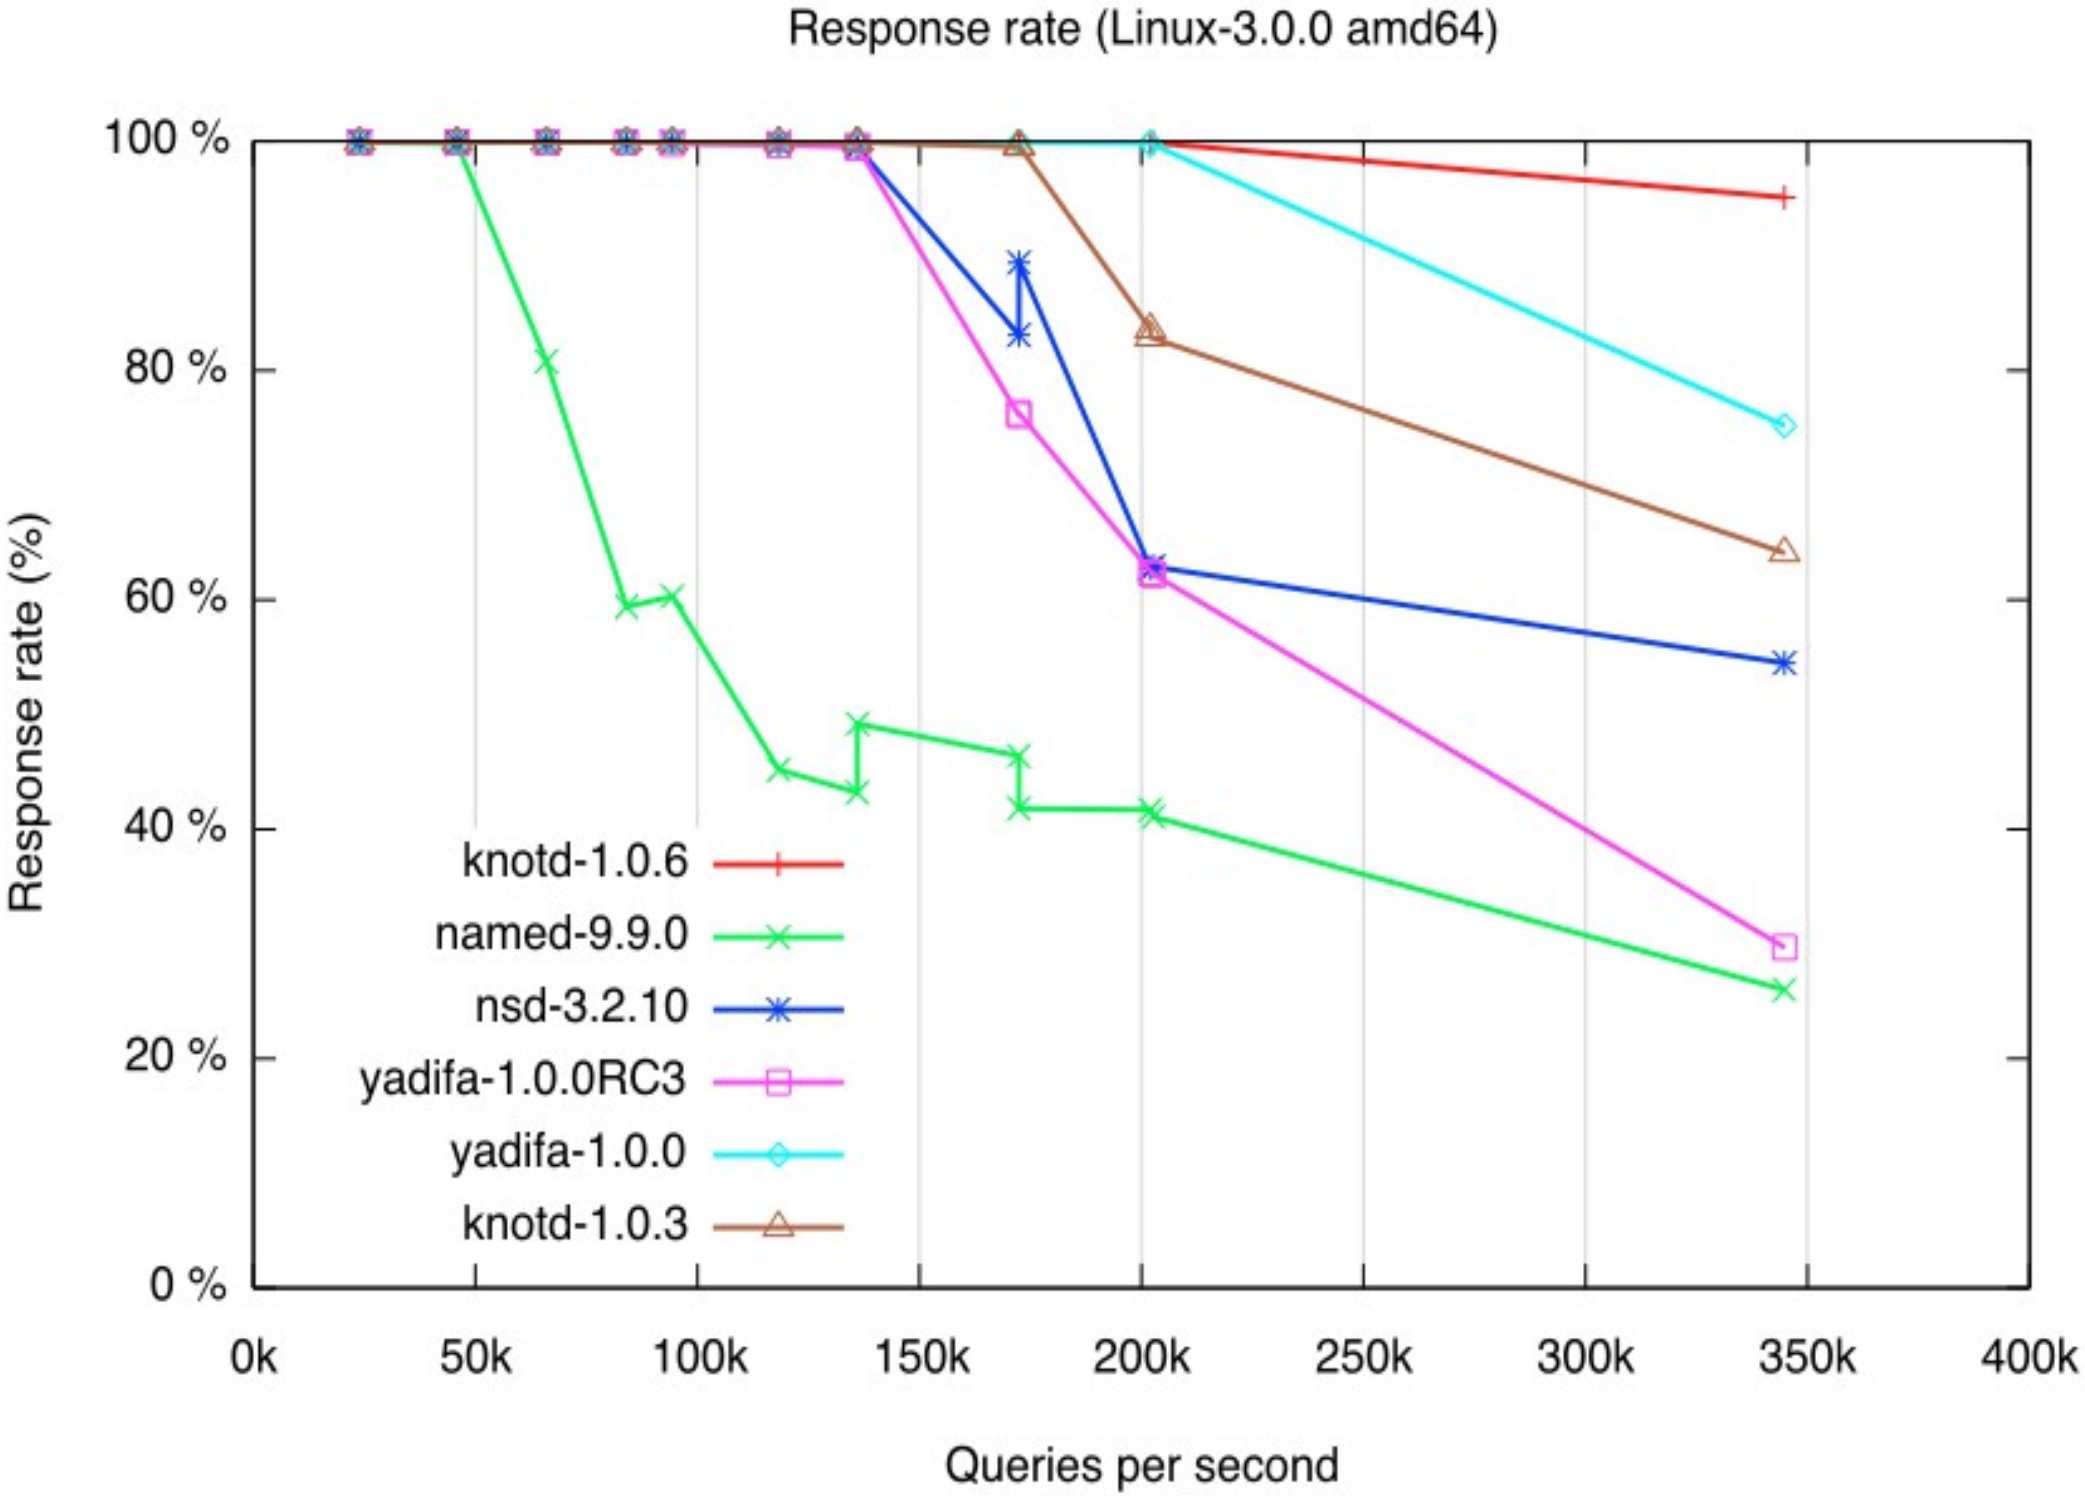
Response rate (Linux-3.0.0 amd64)
100 %
80 %
60 %
40 %
20 %
0 %
Response rate (%)
0k
50k
100k
150k
200k
250k
300k
350k
400k
Queries per second
knotd-1.0.6
named-9.9.0
nsd-3.2.10
yadifa-1.0.0RC3
yadifa-1.0.0
knotd-1.0.3

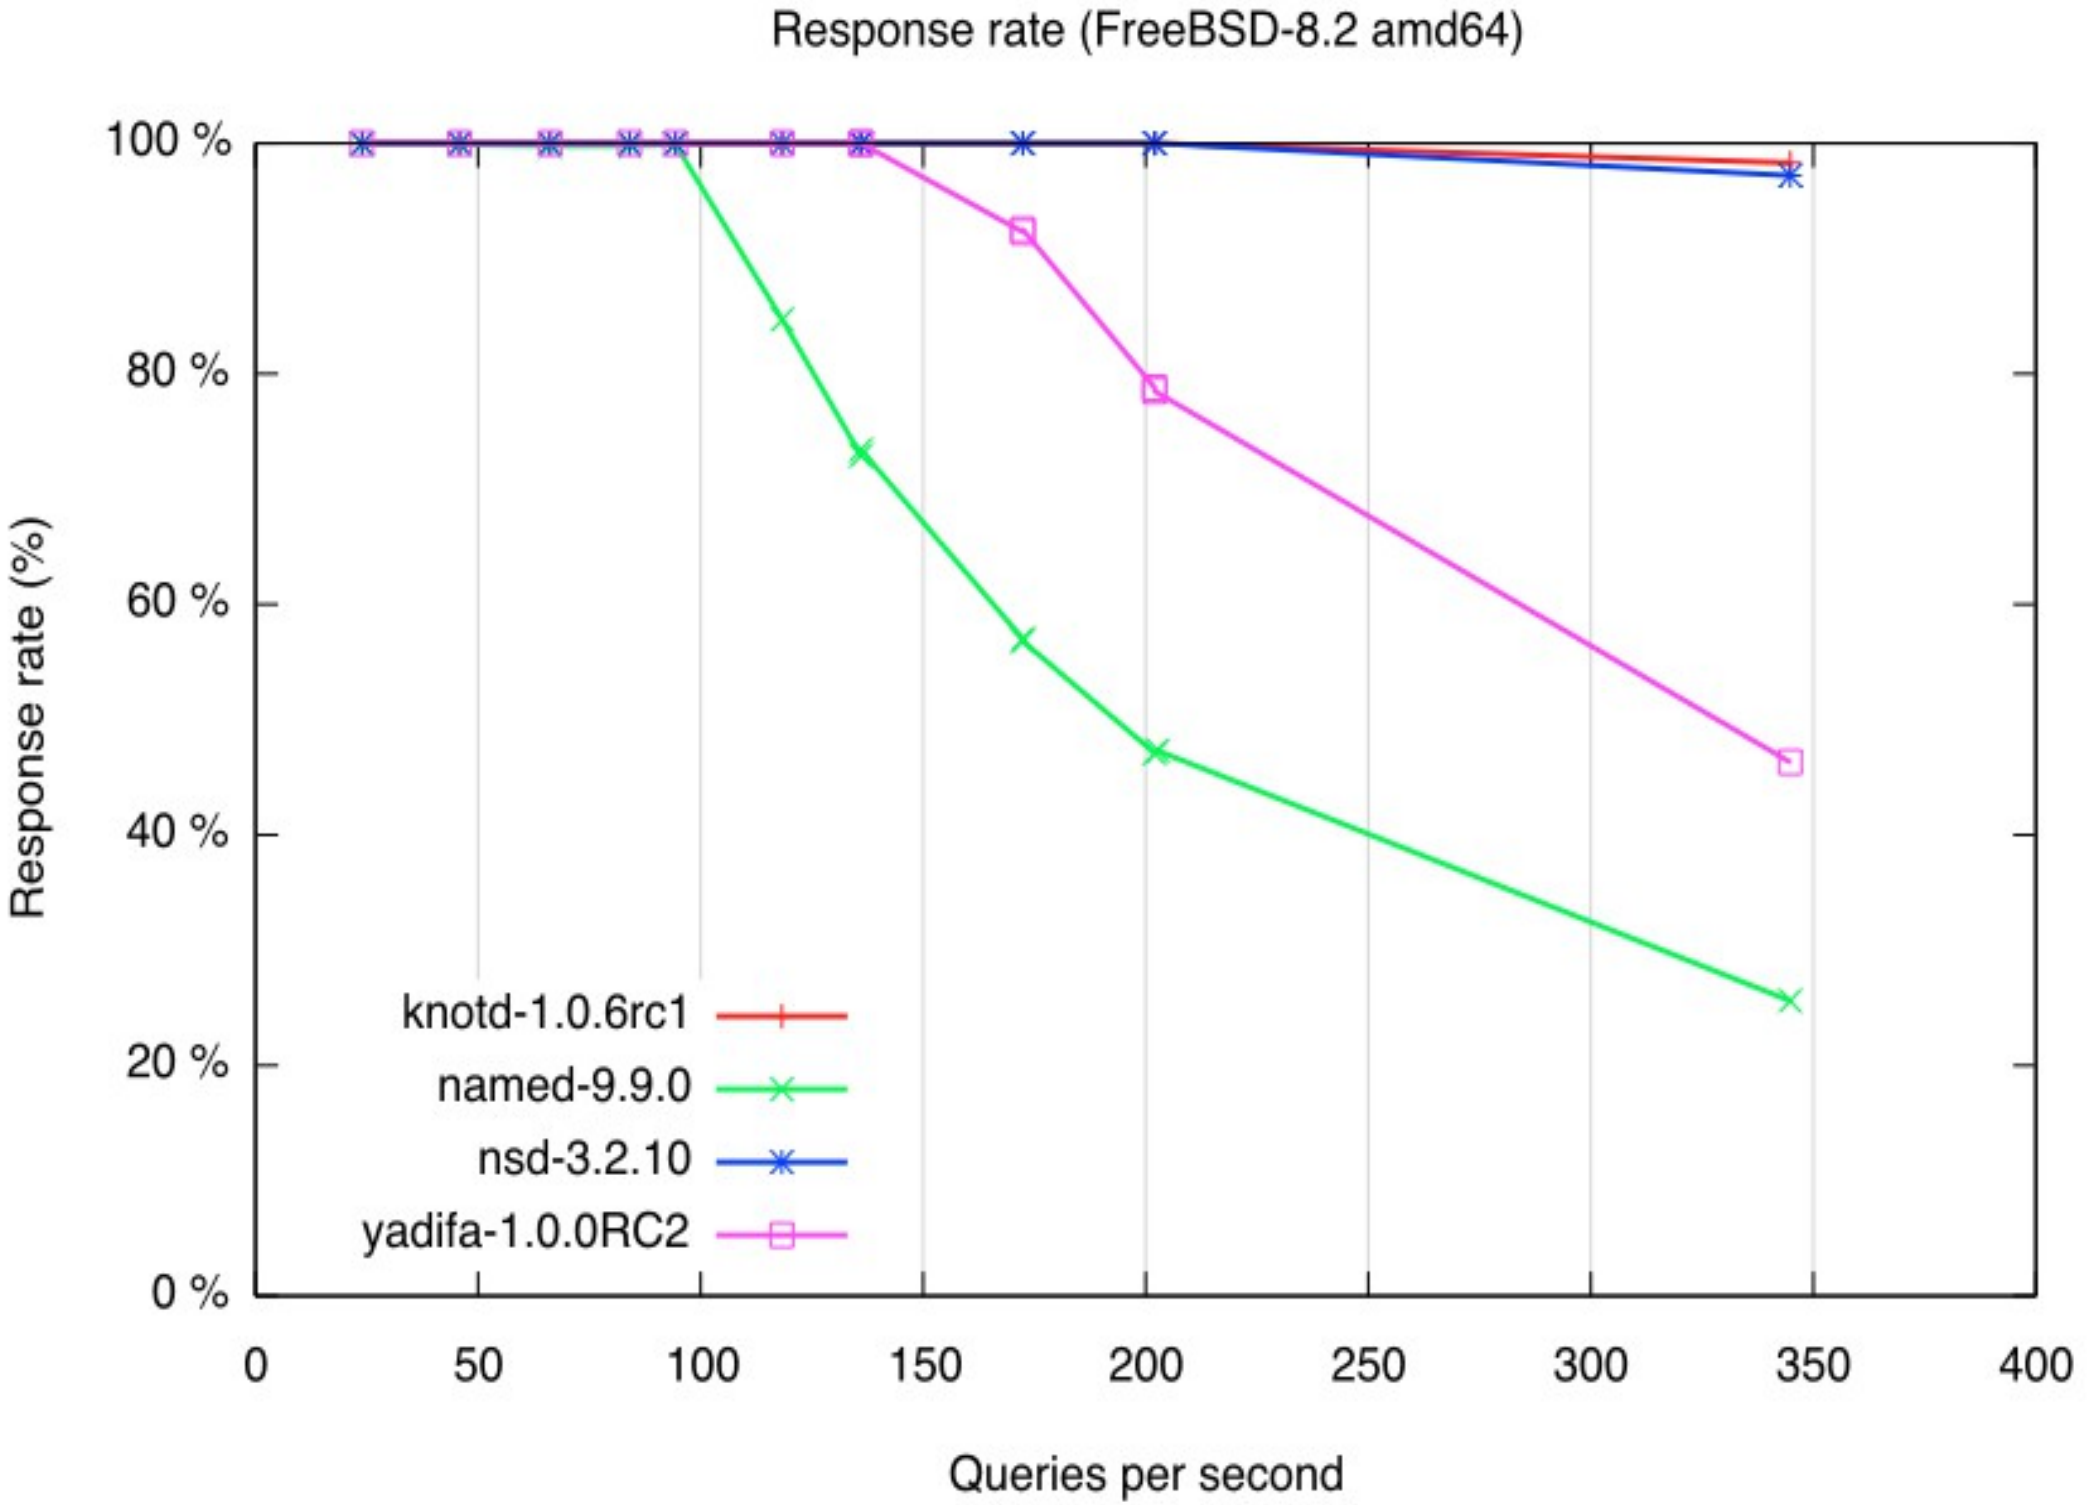
Response rate (FreeBSD-8.2 amd64)
Response rate (%)
100 %
80 %
60 %
40 %
20 %
0 %
0
50
100
150
200
250
300
350
400
Queries per second
knotd-1.0.6rc1
named-9.9.0
nsd-3.2.10
yadifa-1.0.0RC2