It's Knot DNS

CZ.NIC
Ondřej Filip
ondrej.filip@nic.cz
March 21, 2012, DNS-OARC, London, UK

cz nic
cz domain registry

Project (pre-)history

- Started in 2009
  - Small team (Half-man show)
- Speed-up in 2010
  - Two more (half-)people
- September 2011
  - F&F release
  - Small scale deployment

Project history

- November 2011
  - First public release
- 29 February 2012
  - Release 1.0

Project goals

- Open-source authoritative DNS server
  - Alternative to Bind/NSD
- Usable for TLDs (and everybody else)
- Fast
- Portable, modular
- Support for current (useful) standards

Features

- Portable – Linux, *BSDs, Mac OS X
  - Portability now mainly depend on userspace-rcu library
- Pre-compiled zones
- Authoritative-only
- AXFR/IXFR (master and slave)
  - Hopefully we nailed all bugs in IXFR now
- EDNS0
- DNSSEC with NSEC3

Configuration

- Simple curly-bracket based configuration:
  - Interfaces (IPv4/6)
  - Remotes (masters or slaves)
  - Keys (TSIG)
  - Zones (IN class only)
  - Logging (syslog or file-based)
- Runtime reconfiguration
  - Add and remove interfaces
  - Add and remove zones

Configuration example

```
system {
  storage "/var/lib/knot";
}

interfaces {
  lo6 { address ::1@53; }
}

zones {
  example.com {
    file "/etc/knot/example.com.zone";
  }
}

log {
  syslog { any info, warning, error; }
}
```

New features in 1.0

- TSIG and ACLs
  - Subnets now supported in ACL
- Root zone support
- NSID support (RFC5001)
- Automatic zone compile on server start
- Drop privileges after binding to port 53
  - Support for Linux capabilities
- Patched userspace-rcu on {Net,Open}BSD

Design

- C99/GNU99 features
- Object-oriented code
- Modular design
  - Data structure + API
- Mostly lock-free architecture
  - RCU data-synchronization (userspace-rcu library)
- Inspired by BIRD Internet Routing Daemon
  - Also comes from CZ.NIC Labs brewery

Design

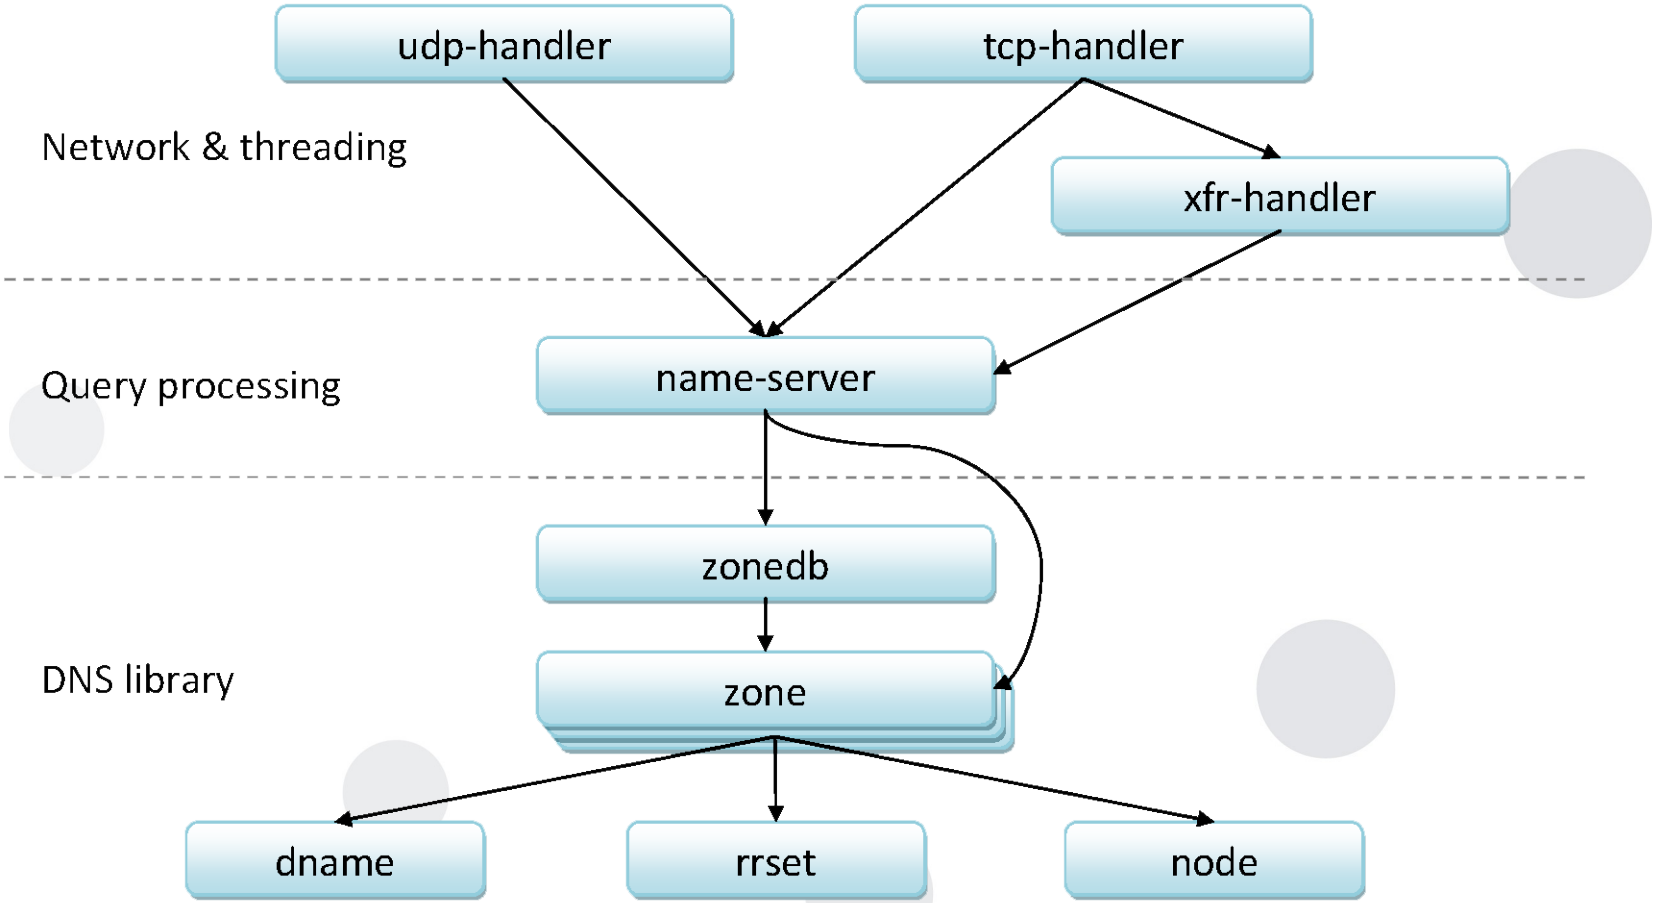
udp-handler
tcp-handler
Network & threading
xfr-handler
Query processing
name-server
zonedb
DNS library
zone
dname
rrset
node

Achieving our goals

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

Linux performance

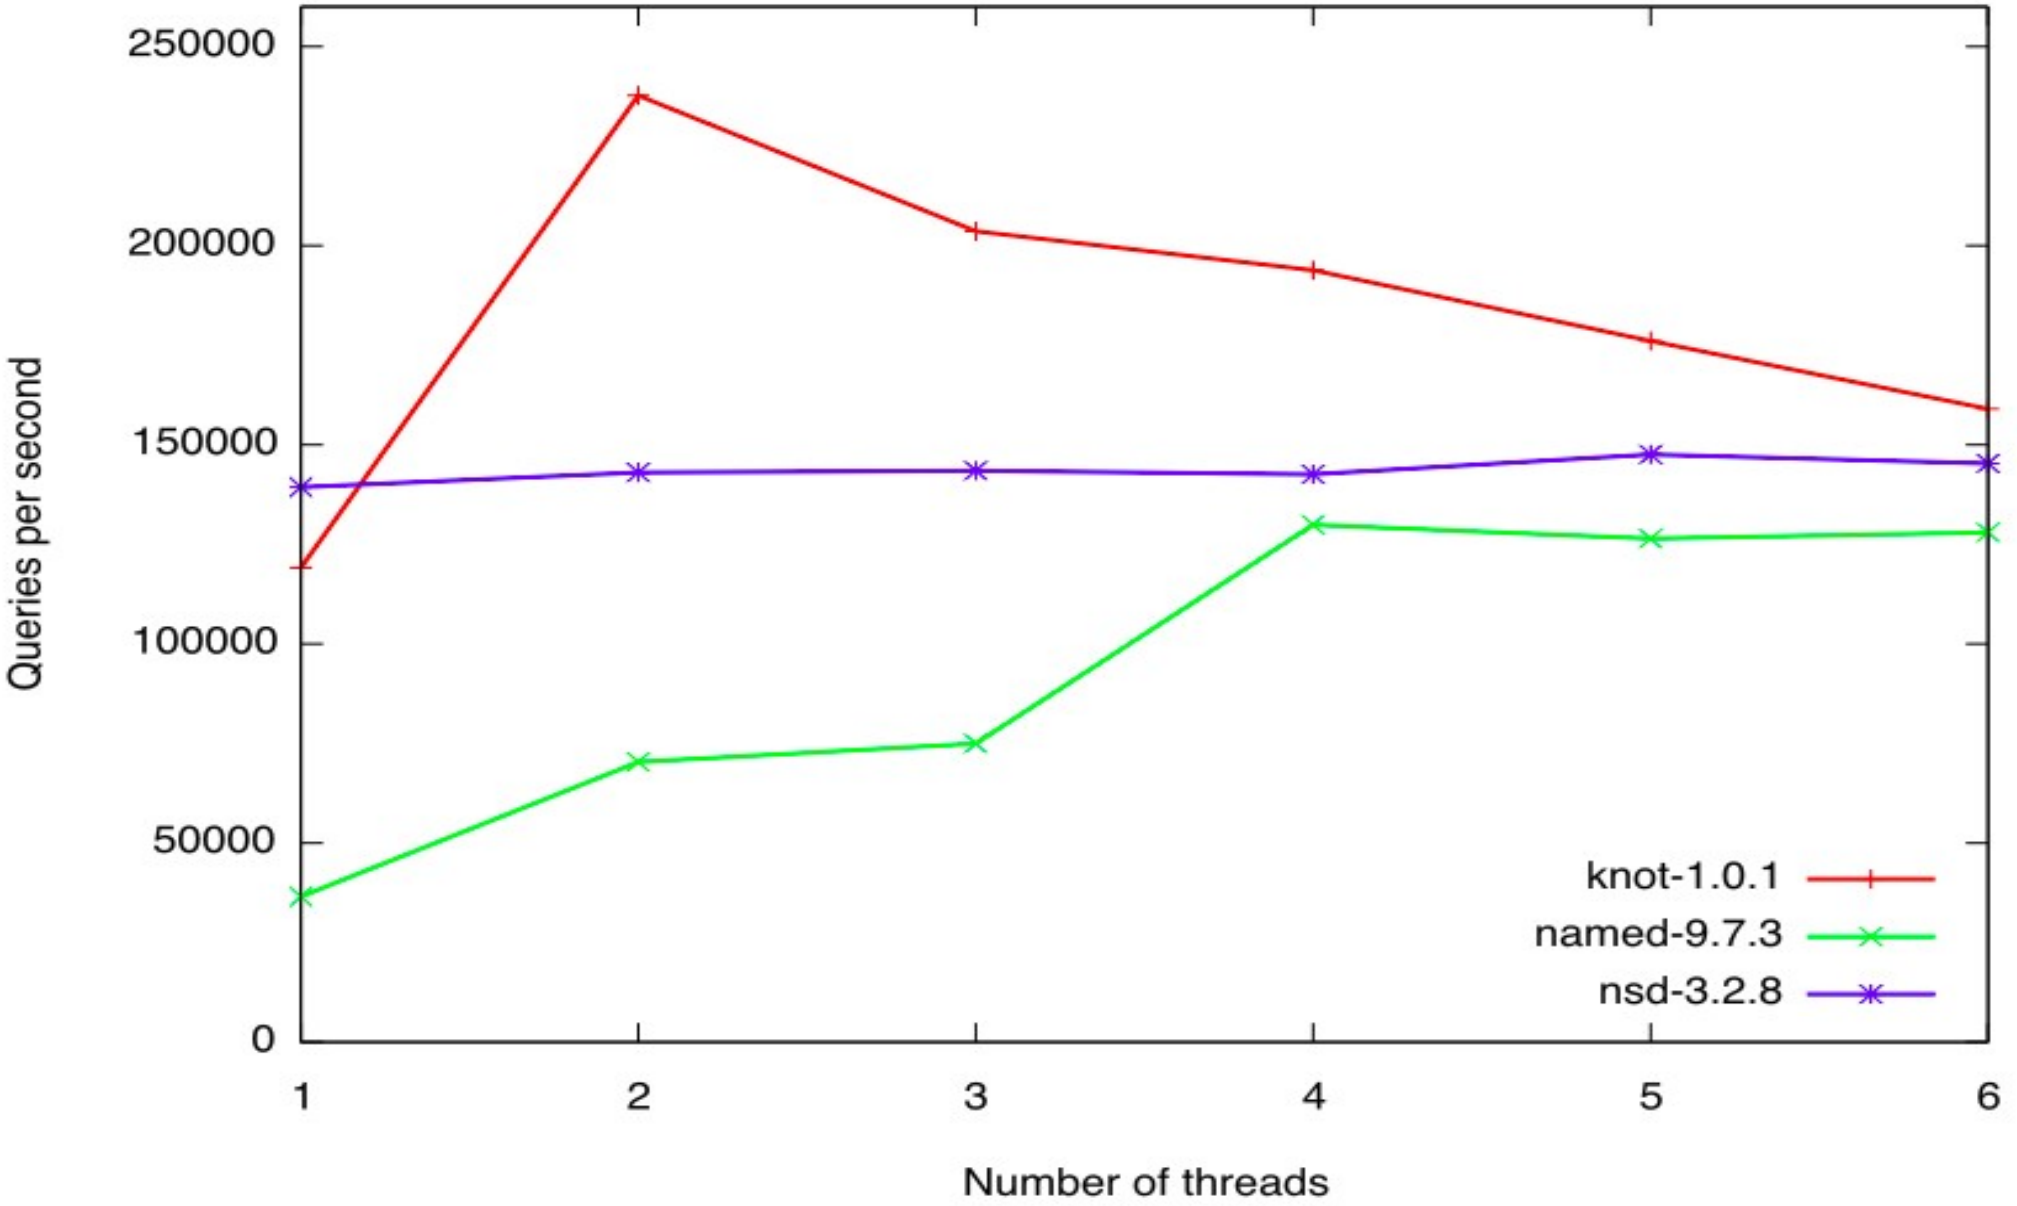

FreeBSD performance

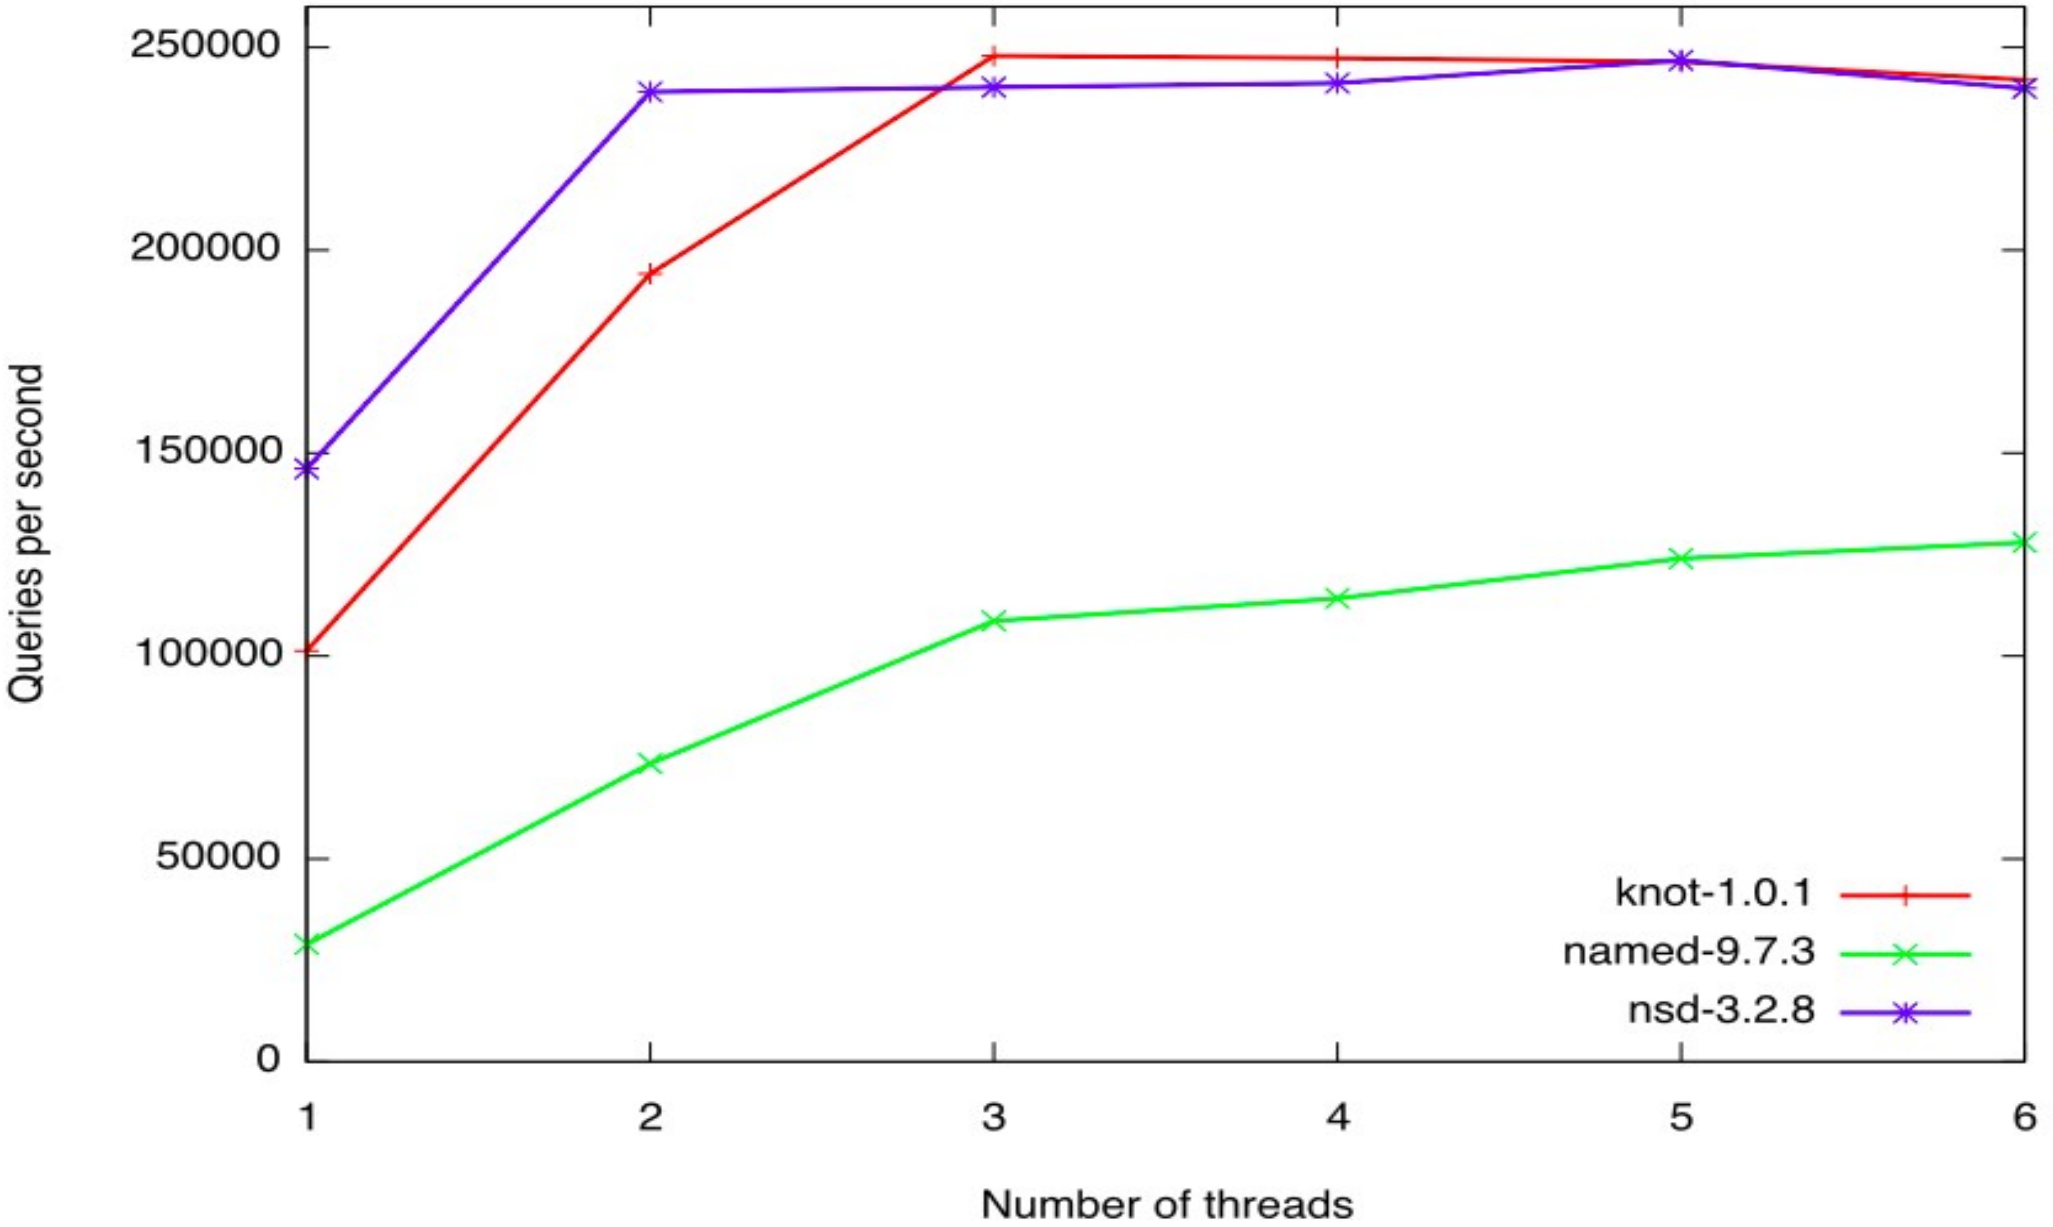

Release plans

- Minor release every 3 months
  - Release early, release often
  - Features in minor releases
- Patch releases as necessary
  - More testing received, more bugs found
  - 1.0.1 already out
  - 1.0.2 in preparations

Release plans[*]

- Knot DNS 1.1 (Q2 2012)
  - Speedup of huge IXFR (10k+ records)
  - Focus on stability and bugfixes
- Knot DNS 1.2 (Q3 2012)
  - Dynamic updates
  - NetConf/DNSCCM support
- Knot DNS 1.3 (Q4 2012)
  - Massive DNS hosting support
  - Slow when loading and serving 10k+ zones

Release plans[*]

- Knot DNS 2.0 (2013)
  - Reduce memory footprint
  - Optimize performance
  - Enhance CLI
- Knot DNS X.Y
  - Your wishes?

* Subject to change
  - Will be published on http://knot-dns.cz/

Summary

PROS

- Performance
- Runtime reconfiguration
- Stability
  - Long-time support from CZ.NIC
- Active Development

CONS

- Higher memory footprint
- Low performance when used with many zones
- Need more testing

Resources

- Knot DNS
  - http://www.knot-dns.cz/

    (Language switcher at top of the page, WIP)
- Issue tracking and source code
  - http://git.nic.cz/redmine/
  - git://git.nic.cz/knot-dns
- Mailing list

  knot-dns-users@lists.nic.cz

Questions?

cz nic

cz domain registry