

It's Knot DNS!



CZ.NIC, z.s.p.o.

Ľuboš Slovák, Marek Vavruša

lubos.slovak@nic.cz, marek.vavrusa@nic.cz

5.11.2011


Project goals

- Open-source DNS server
 - Alternative to Bind/NSD
- Fast(est), feature-rich & on-the-fly reconfiguration
- Usable for TLDs
- Portable, modular
- Support all current (useful) standards




Project history

- Started in 2009
 - Very small development team
- Development boosted in 2010
 - Two more people
- September 2011
 - Friends & Family Release
 - Small scale deployment
- November 2011
 - Public release - RIPE63



Features

- Portable – Linux, *BSDs, Mac OS X
- DNS standards implemented
 - Authoritative only
 - AXFR/IXFR (both master and slave)
 - ACLs
 - EDNS0
 - DNSSEC + NSEC3
 - All specified RR types
 - Unknown RR types



Features


- Simple C-like configuration file
 - Interfaces (IPv4 or IPv6)
 - Remotes (masters or slaves)
 - Zones
 - Logging (syslog or file-based)
- Runtime reconfiguration
 - Add and remove interfaces
 - Add and remove zones
- Zone compilation
- Simple control interface (knotc)



Design

- Object-oriented code
- Modular design
 - Structure + API
- Mostly lock-free architecture
 - RCU data-synchronization (liburcu)
- Inspired by BIRD Internet Routing Daemon
 - Other CZ.NIC Labs product

Design



Achieving our goals: Network

- Possible bottleneck
- Hardware is important
 - Choice of network card & drivers
- Sockets
 - I/O buffer size (setsockopt)
 - Scatter / gather I/O
 - Own wrapper for managing TCP connections
 - Portability

Achieving our goals: Threading

- Dynamic threads implementation
- Minimizing amount of syscalls per query
 - Blocking vs. non-blocking sockets
 - Linux – `recvmsg ()`

Achieving our goals: query processing

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



Live demo

Benchmarks

- Setup

- 4-core Intel Xeon X3430, 2.40 GHz, 2 GB RAM
- Linux 2.6.38-11, x86_64
- FreeBSD 8.2, x86_64


- Tools

- dnssperf (by Nominum)


- Data

- Whole CZ zone with NSEC3

Response rate - Linux



Response rate – FreeBSD



Testing

- Unit tests
- Overall server testing
 - Evaluating results using different implementation
 - GoDNS
- Regression testing
 - Comparing to BIND

Knot DNS at CZ.NIC

- Deployed at some less-important SLDs
 - e.g. knot-dns.cz
- Deployment plan for .CZ
 - Before end of 2011
 - Only selected anycast nodes

(Near) Future Plans

- Feature-freeze
 - TSIG
 - Dynamic updates
 - NSID
 - Root zone support
- Testing, debugging, fixing bugs, testing, ...



Future Plans

- Even better performance
- Reduce memory footprint
 - Without reducing performance
- Richer CLI
 - Remote access
 - Automatic zone compilation

Summary

PROS

- Performance
- Runtime reconfiguration
- Developed in the DNS (TLD) operator community
- Active development

CONS

- Not-yet feature complete
- Higher memory footprint
- Possible bugs (beta)

Conclusion

- New high-performance authoritative server
- Beta-version
 - Not for production use yet
- Tarball or packages
 - deb (Debian and Ubuntu)
 - rpm (Fedora)
- Feedback and testing welcome

More resources

- Knot DNS

<http://www.knot-dns.cz/>

- Development site (issue tracking)

<https://git.nic.cz/redmine/projects/knot-dns>

- Git repository

<git://git.nic.cz/knot-dns>

- Mailing list

knot-dns-users@lists.nic.cz

Questions?



<http://www.knot-dns.cz/>

knot-dns-users@lists.nic.cz

