



NLnet Labs Product Roadmap

OARC 38 Vendor Panel

31 July 2022





NSD: Recent features & plans

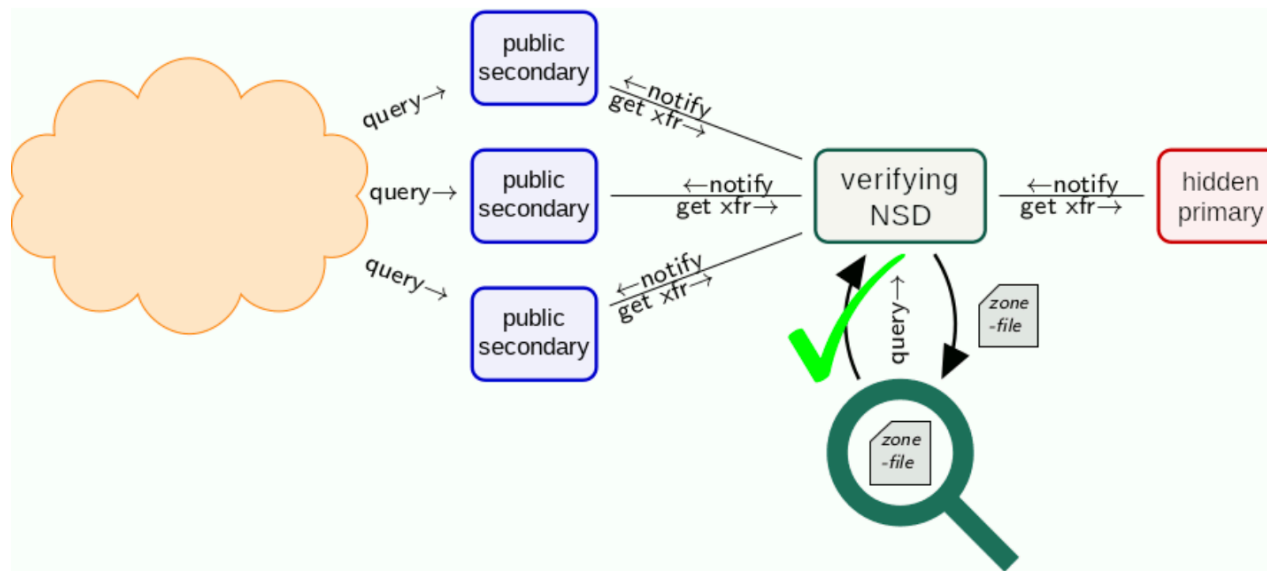


NSD 4 Recent Results

- Strategic development
 - Move NSD 4 towards/closer to provisioning systems
 - introduce features to deploy NSD as master server
 - integration with backend systems
 - Performance and stability are prime focus
- Keeping up with new IETF standards
 - Extended DNS Errors, SVCB/HTTPS RR types, interoperable DNS cookies support, DNS zone transfer over TLS

NSD 4 Features in Past Months

- IXFR-out
 - best effort, not arbitrary delta's but most recent
- Zone verification (formerly known as CreDNS)



NSD Roadmap for 2022-2023

- Performance features (under development)
 - Redesign zone parsing and loading
 - Alternative data structure to red-black and radix tree, e.g. adaptive radix tree
 - XDP (EXpress Data Path) support
- Catalog zones (soon WGLC)
- DoQ (DNS-over-QUIC)
 - DoT already supported
 - unilateral probing draft in DPRIVE good candidate to reach the finish line



Unbound: Recent features & plans

Unbound Strategic Developments

- Unbound resiliency
 - serve stale, operational behaviour following CNAMEs, timeout and retry strategies
- Enterprise
 - RPZ
 - views per interface, ACLs per interface
- Cloud providers/CDN
 - upstream server selection optimisation/tuneable

Unbound Recent Results

- ZONEMD
- SVCB/HTTPS support (types and handling)
- TCP/TLS stream reuse: performing several queries over the same TCP or TLS channel
- Extended DNS Errors

Unbound Roadmap for 2022-2023

- Under development
 - DoQ client side
 - Proxy protocol version 2, client side
 - ACL per interface
 - DNS cookies (upstream)
 - DNS Error reporting (happy path during IETF 114 Hackathon)

Unbound Roadmap for 2022-2023

- Future work
 - DoQ upstream
 - Unilateral probing for encrypted DNS to upstream
 - Prometheus metrics support?
 - DNS views (generic)



Knot DNS news



**KNOT
DNS**

Libor Peltan • libor.peltan@nic.cz • 2022-07-31

Knot DNS versions

1

- Knot DNS 3.0 stable
 - Sep 2020
- Knot DNS 3.1 current
 - Aug 2021
- Knot DNS 3.2 soon
 - Aug/Sep 2022

Use our repositories
<https://www.knot-dns.cz/download/>
Distribution's get out-of-date



XDP: performance and DoS resilience

2

- XDP-UDP stable, good, used
- XDP-TCP improved 3.1 → 3.2
 - DoS resilience for public-facing servers
 - Not suitable yet for AXFR-out
- XDP-QUIC upcoming release, first version
 - Including kdig, kxdpgun



Many-zones AXFR/IXFR improvements

3

- Retry lost NOTIFY
 - Avoid redundant NOTIFY
- Reusing TCP connections
- Simplified configuration
 - Automatic ACLs
 - Grouping of remotes



Catalog zones

4

- Improvements 3.0 → 3.1 → 3.2
- Consumption & generation incl. Group property
- <https://datatracker.ietf.org/doc/html/draft-ietf-dnsop-dns-catalog-zones>
 - In cooperation with all other vendors
 - Incl. IETF 113 Hackathon



- Memory consumption processing big changesets
- Multi-signer scheme helpers
- DBus hooks for substantial events in server
- EDNS Expire option (RFC 7314)

- ...and much more





Knot Resolver news



**KNOT
RESOLVER**

Vladimir Cunat • vladimir.cunat@nic.cz • 2022-07-31

Nameserver choice

- Which IP address to ask? Resolve other NS names?
- Old algorithm: too magical, sometimes bad behavior
- 5.3.0: new implementation
- Better latency on cache miss, fewer packets
- Sharing NS stats through LMDB cache



Assertions

- Dilemma – if internal inconsistency is detected:
 - Abort: coredumps help fixing these bugs
 - Recover: minimize disturbance of the DNS service
- Do both: fork
 - child aborts
 - parent recovers
- Since 5.4.0



Logging

- Metadata
 - < 5.4.0: just plain text, now syslog + systemd API
 - Logging level, usage e.g. `journalctl -p err`
 - With systemd: also precise position in source code
- Groups
 - Each log line marked with group, e.g. `[cache]`
 - Activating `debug` level for each group separately



Miscellaneous in 5.5.0

- Extended EDNS errors, RFC 8914
- ZONEMD: validation of prefilled root zone
 - code inactive by default, waiting for root
- PROXYv2 support on client side



An update on PowerDNS

Peter van Dijk

PowerDNS Senior Engineer

OARC38 Vendor Panel

General

- time_t
- CentOS 8 -> Oracle Linux 8
- FALCON-512
- embedded deployment

PowerDNS Recursor

4.5 (May 2021)

- 8198: Aggressive use of DNSSEC-Validated Cache
- EDNS0 padding (RFC 7830) towards clients
- DNSSEC validation enabled by default
- 8914 extended errors
- refetch
- outgoing TCP Fast Open
- non-resolving nameserver cache (minor TsuNAME mitigation)

PowerDNS Recursor

4.6 (December 2021)

- flush-on-NOTIFY
- DoT to auth (manual config)
- TCP/DoT connection reuse
- zone-to-cache (XFR, web, file)

PowerDNS Recursor

4.7 (May 2022)

- free additional records for clients
- updated qname minimization (RFC9156)
- active fetching of IPv6 glue
- unilateral DoT probing
- ZONEMD validation for zone-to-cache

PowerDNS Authoritative Server

4.5 (July 2021)

- zone name cache
- AXFR-in queue priority ordering

PowerDNS Authoritative Server

4.6 (January 2022)

- incoming PROXYv2 headers
- EDNS cookies
- improved default NSEC3 parameters

PowerDNS Authoritative Server

4.7 (soon)

- catalog zones!

dnsmist

1.6 (May 2021)

- OoO
- PROXYv2
- lockless balancing policies in Lua
- cookie-blind packetcache


dnsmist

1.7 (January 2022)

- outgoing DoT+DoH
- TC responses via XDP
- SVCB/HTTPS
- lockless custom actions in Lua

BIND 9 Update, OARC 38

Vicky Risk, vicky@isc.org



← Tweet



Tomasz Łakomy ⚡ cloudash.dev 🇺🇦

@tlakomy



Fixing tech debt in an enterprise codebase



5:26 PM · Jul 13, 2022 · Twitter for iPhone



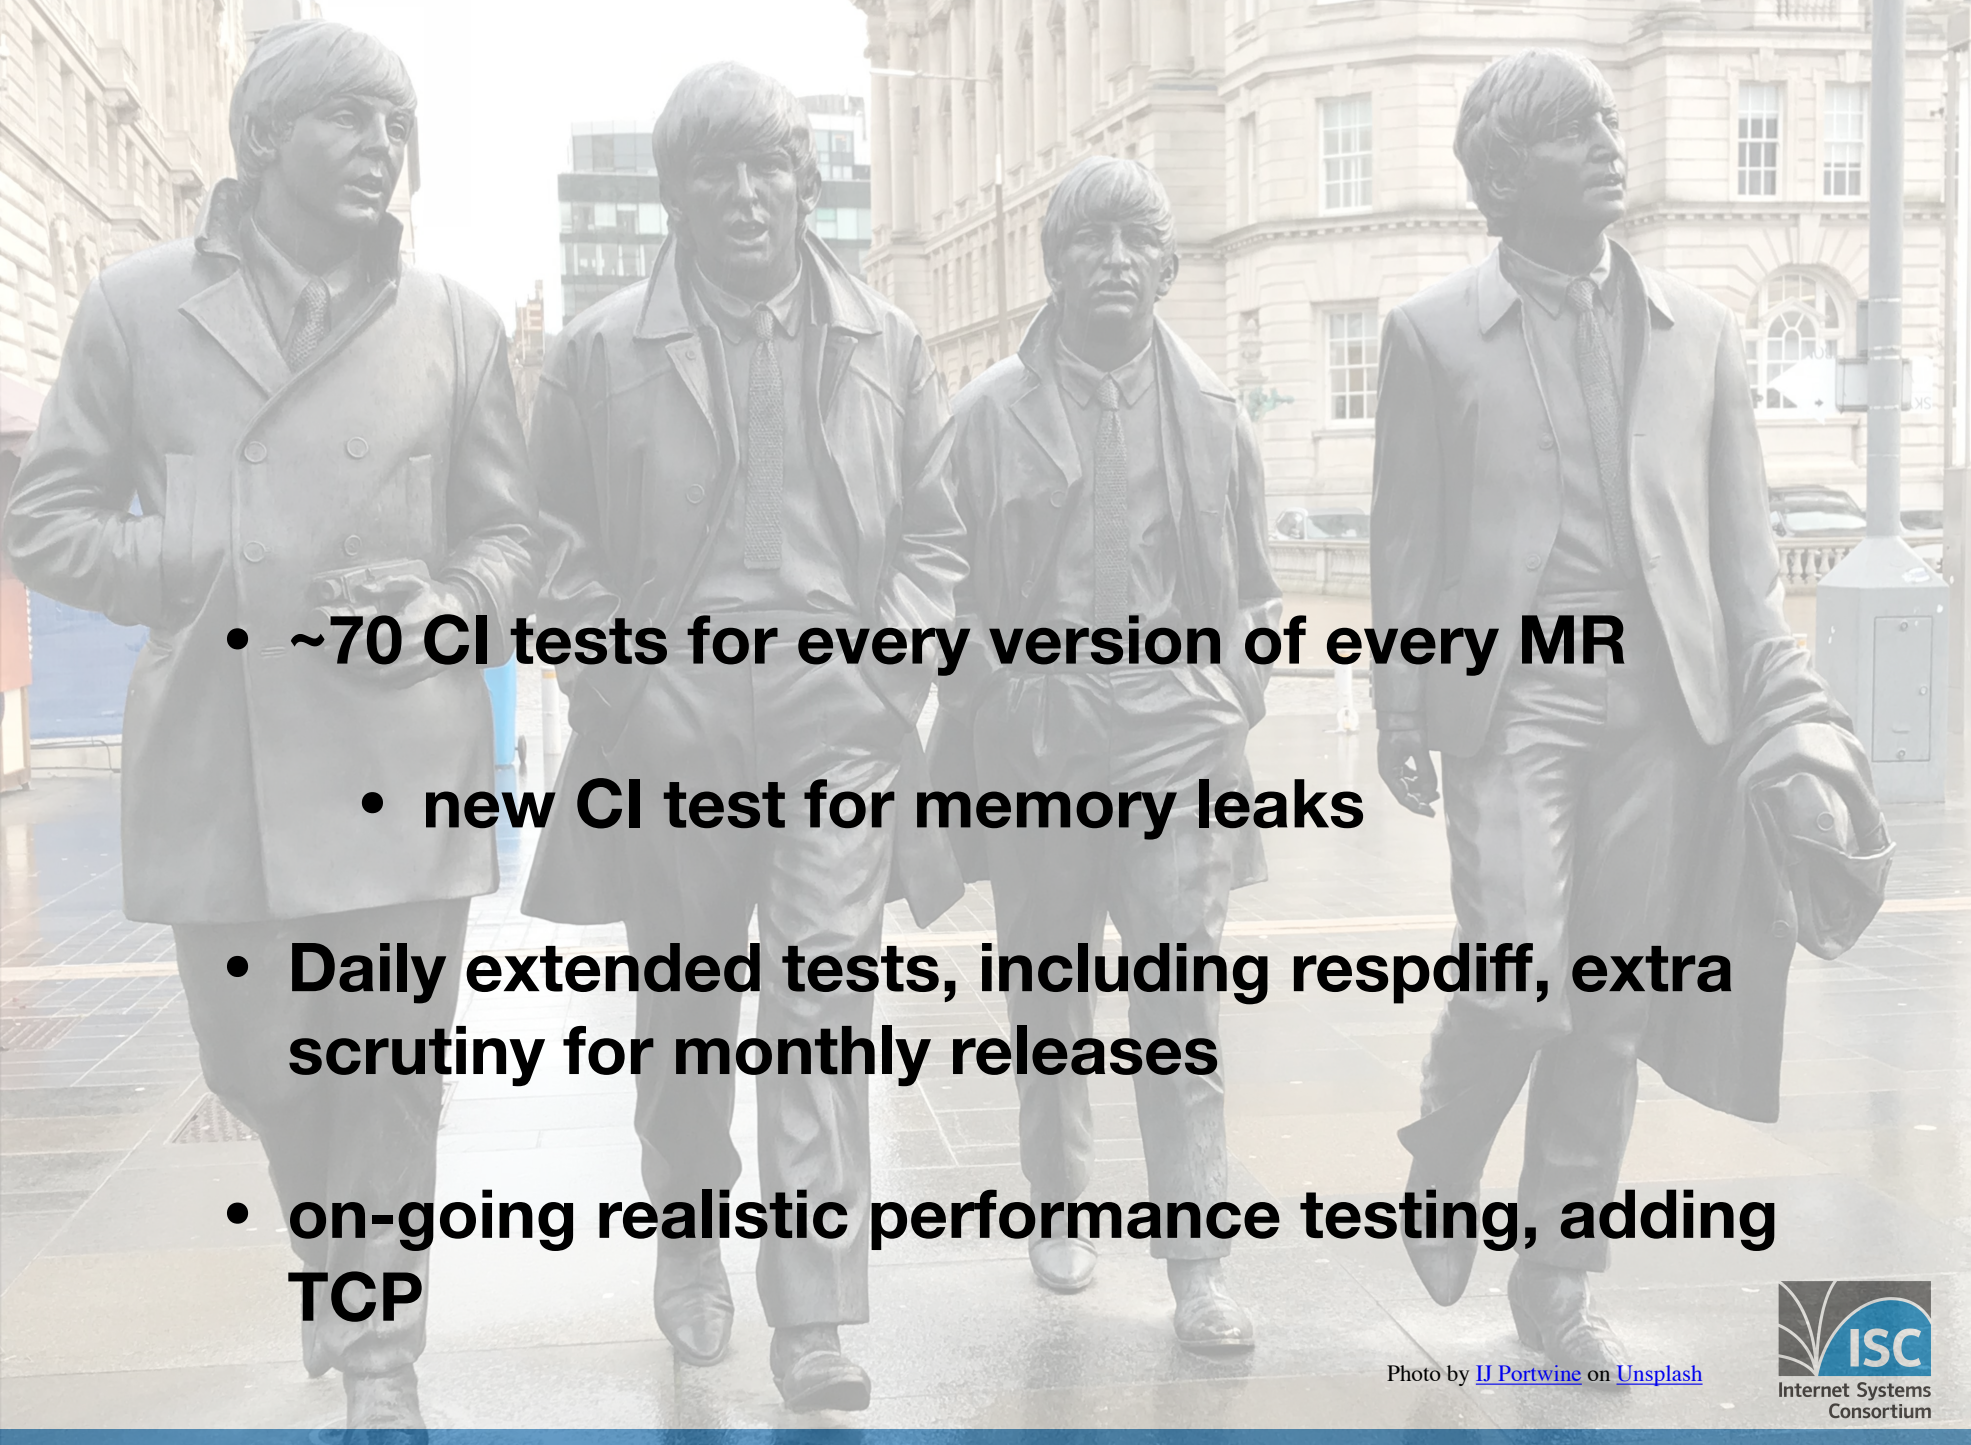
- 
- **~70 CI tests for every version of every MR**
 - **new CI test for memory leaks**
 - **Daily extended tests, including respdiff, extra scrutiny for monthly releases**
 - **on-going realistic performance testing, adding TCP**

Photo by [IJ Portwine](#) on [Unsplash](#)

Testing - respdiff

we're comparing what **NAMED** returns for a predefined set of queries against what Knot Resolver (5.5.1), Unbound (1.13.1), and PowerDNS Recursor (4.5.9) return

if the *ratio* of discrepancies exceeds a preset threshold (0.5%), the job fails (we investigate a possible BIND error)

note, however, how this is calculated: a difference only counts if **all** other resolvers return the same response and **NAMED** returns something different

named disagrees

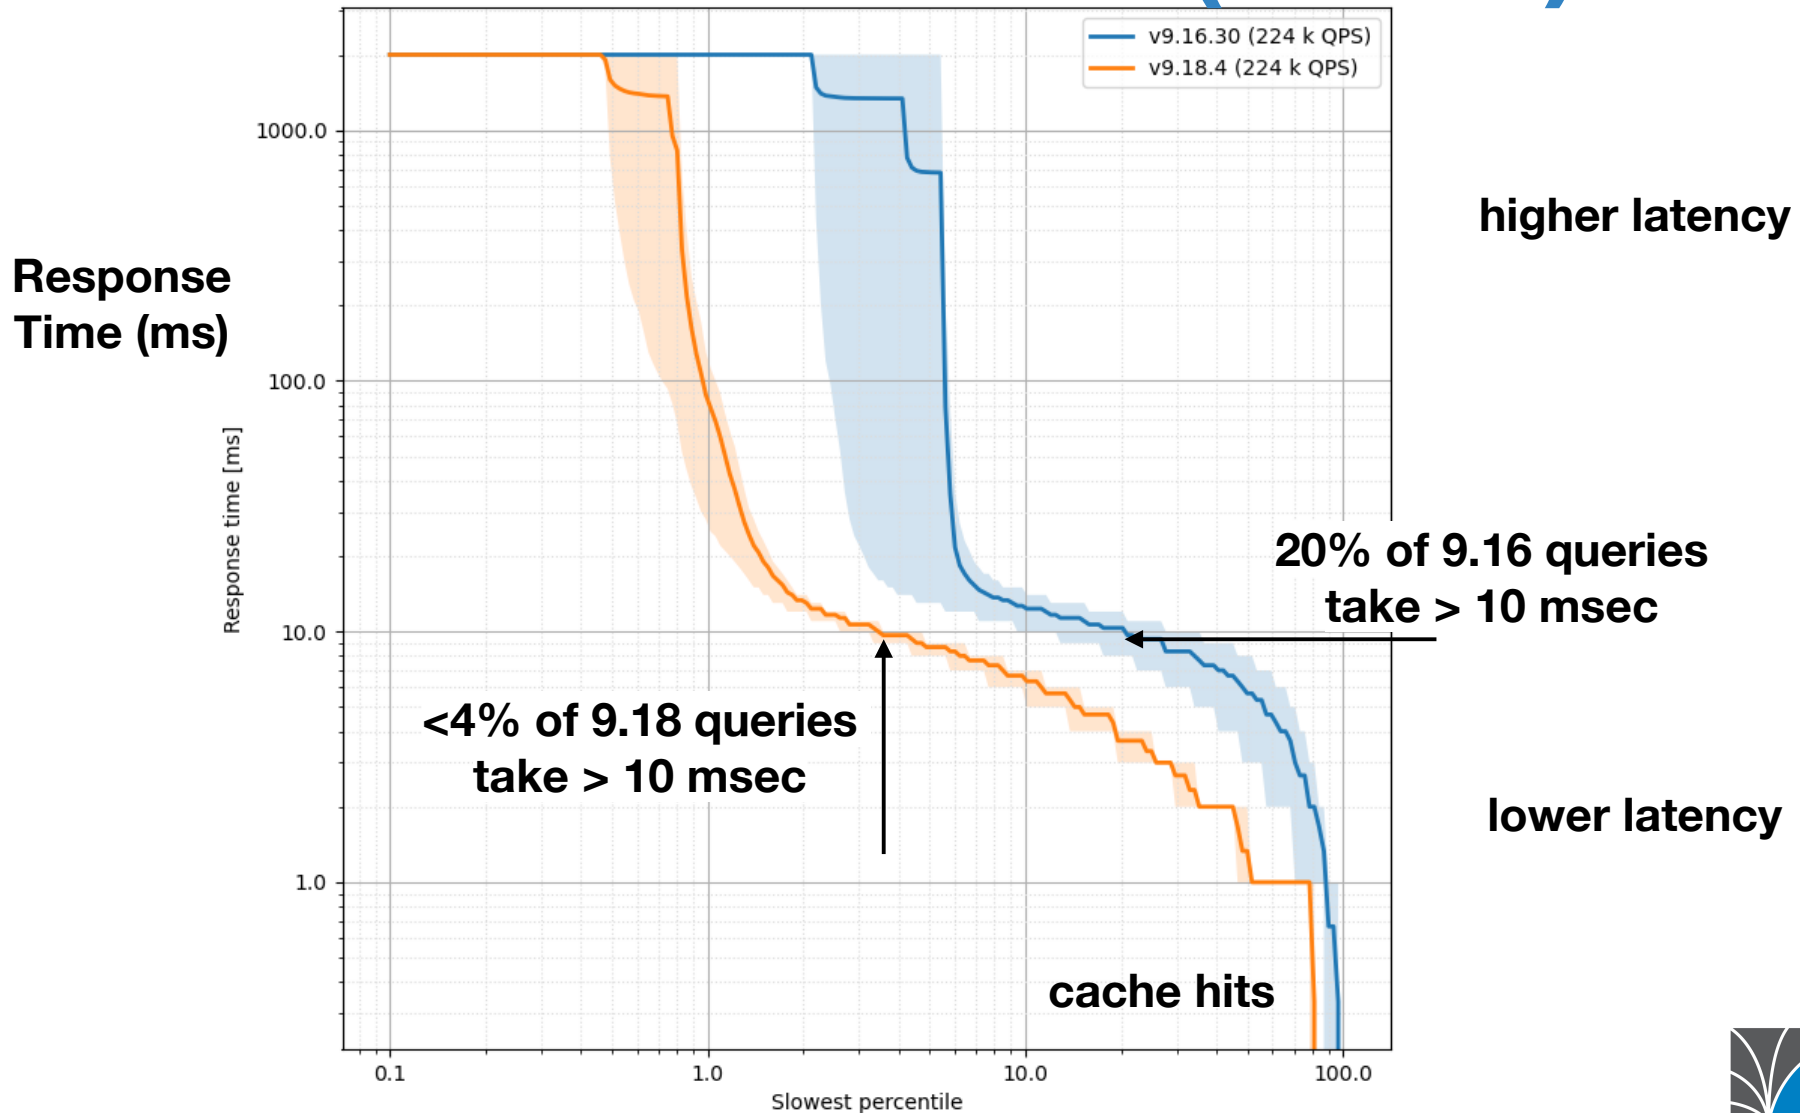
0.34%

Target disagreements between the tested version and the reference one (Knot Resolver 5.5.1, Unbound 1.13.1, and PowerDNS Recursor 4.5.9) comprise **0.34%** of not ignored answers; of these, 38.11% are timeout disagreements, which can be attributed to network issues*.

* Network differences occur even between separate runs with the same software, due to the natural variability of the live Internet.

9.16 vs 9.18 (UDP)

UDP, 24x load



see <https://www.isc.org/blogs/bind-resolver-performance-july-2021/> for test bed description



Replacing RBTDB

Goals

1. Code simplification
2. reduce blocking on updates
3. not slower
4. not more memory

Plan

- Adapt qp-trie, invented by Tony Finch in 2015
- used by Knot DNS since 2016, experiments using NSD 2020-2021
- more complete multithreading, multi-version concurrency
- test in isolation, replace rbtodb in stages

Photo by [Jan Antonin Kolar](#) on [Unsplash](#)

qp-trie Status

**in early testing before merge
into main (dev) branch**

**single-threaded code solid,
multithreaded in progress**

very preliminary benchmark	RBT	qp-trie
time to load 1 M domain names	1.0 seconds	0.7 seconds
memory consumed	113.3 MiB	45.4 MiB

research code, blog articles, and notes: <https://dotat.at/prog/qp/>

work-in-progress branch (unstable, probably broken): <https://gitlab.isc.org/isc-projects/bind9/-/commits/fanf-qp-import>



other WIP

- moar Extended Errors
- catalog zones update to the 06 draft
- sponsoring an OpenSSL 3.0 PKCS #11 provider engine
- refactoring: “stream DNS” for TLS & TCP
- ARM update with Ron Aitchison (ProDNS and BIND author)

Photo by [Mike Kenneally](#) on [Unsplash](#)

Linking, tagging in the ARM

allow-notify

Grammar: `allow-notify { <address_match_element>; ... };`

Blocks: options, view, zone (mirror, secondary)

Tags: transfer

Defines an `address_match_list` that is allowed to send NOTIFY messages in addition to addresses defined in the `primaries` option.

This ACL specifies which hosts may send NOTIFY messages for zones for which it is acting as a secondary server. This (i.e., `type secondary` or `slave`).

If this option is set in `view` or `options`, it is globally applicable. In a `zone` statement, the global value is overridden.

If not specified, the default is to process NOTIFY messages for the zone. `allow-notify` can be used to expand the

Links

8.3.2. Transfer Tag Statements

Statement	Description
<code>allow-notify</code>	Defines an <code>address_match_list</code> that is allowed to send NOTIFY messages for the zone, in addition to addresses defined in the <code>primaries</code> option for the zone.
<code>allow-transfer</code>	Defines an <code>address_match_list</code> of hosts that are allowed to transfer the zone information from this server.
<code>allow-update</code>	Defines an <code>address_match_list</code> of hosts that are allowed to submit dynamic updates for primary zones.
<code>allow-update-forwarding</code>	Defines an <code>address_match_list</code> of hosts that are allowed to submit dynamic updates to a secondary server for transmission to a primary.
<code>also-notify</code>	Defines one or more hosts that are sent NOTIFY messages when zone changes occur.
<code>alt-transfer-source</code>	Defines alternate local IPv4 address(es) to be used by the server for inbound zone transfers, if the address(es) defined by <code>transfer-source</code> fail and <code>use-alt-transfer-source</code> is enabled.
<code>alt-transfer-source-v6</code>	Defines alternate local IPv6 address(es) to be used by the server for inbound zone transfers.
<code>ixfr-from-differences</code>	Controls how IXFR transfers are calculated.
<code>max-journal-size</code>	Controls the size of journal files.

Parent vs Child



Photo by [Kelli McClintock](#) on [Unsplash](#)

We are *debating* changing BIND from child-centric to parent-centric.

Your comments are welcome at:

<https://gitlab.isc.org/isc-projects/bind9/-/issues/3311>

Current Stable - 9.18

Now on a 2-year development cycle. 9.18 released Jan 2022.

- Better recursive performance
- Better memory usage, reduced fragmentation with jemalloc
- TLS security - DoH, DoT & XoT, dig +tls
- DNSSEC KASP
- OpenSSL 3.0
- Extended errors

Links

- Tony's blog (qp-trie): <https://dotat.at/prog/qp/>
- Parent vs child discussion: <https://gitlab.isc.org/isc-projects/bind9/-/issues/3311>
- Stream DNS issue: <https://gitlab.isc.org/isc-projects/bind9/-/issues/3374>
- example res-diff output: <https://gitlab.isc.org/isc-projects/bind9/-/jobs/2635397>