BIND 9 Update, OARC 38

Vicky Risk, vicky@isc.org

Internet Systems Consortium

← **Tweet**

Tomasz Łakomy cloudash.dev
@tlakomy

Fixing tech debt in an enterprise codebase

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

Internet Systems Consortium

Expanded QA team

- ~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

Internet Systems Consortium

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

Internet Systems Consortium

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.

ISC
Internet Systems Consortium

9.16 vs 9.18 (UDP)

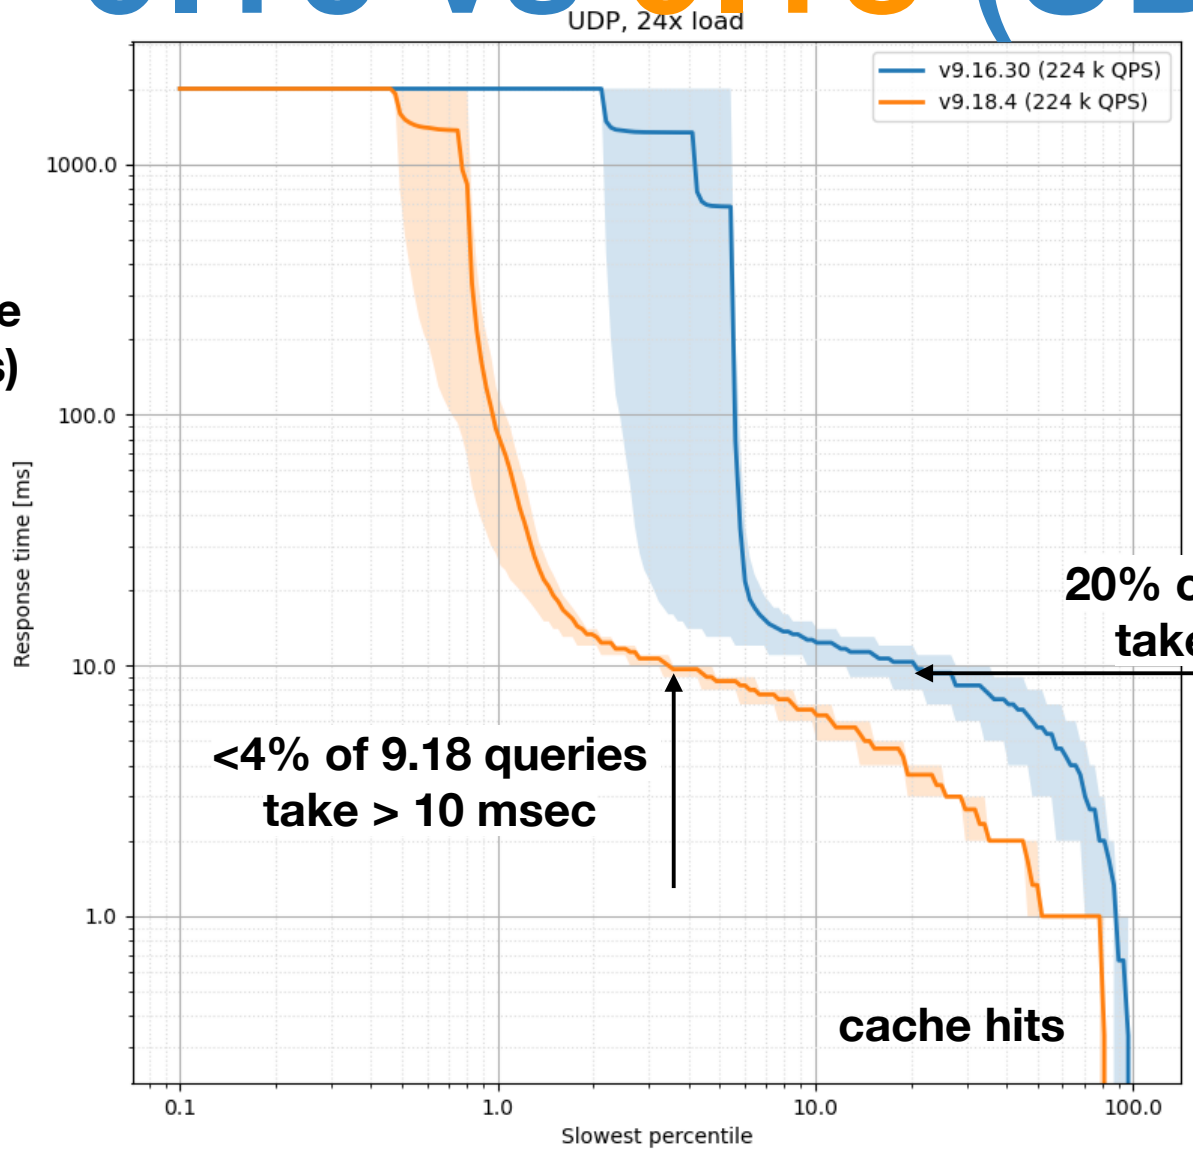

Response Time (ms)

higher latency

20% of 9.16 queries take > 10 msec

<4% of 9.18 queries take > 10 msec

lower latency

cache hits

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

Internet Systems Consortium

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 rbtdb in stages

Photo by Jan Antonin Kolar on Unsplash

ISC Internet Systems Consortium

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

ISC Internet Systems Consortium

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

Internet Systems Consortium

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
addition to addresses defined in the **`primaries`** option

This ACL specifies which hosts may send NOTIFY mes
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 a
`zone` statement, the global value is overridden.

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

Links

8.3.2. Transfer Tag Statements

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

Internet Systems Consortium

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

Internet Systems Consortium