

AGM October 2024

Engineering and Software Report

Matt Pounsett, Michael Baer, Jerry Lundström

AGM 2024 - Engineering Update

Matt Pounsett, Michael Baer

OARC Services Overview

Data Archiving

OARC maintains a large store of multiple datasets:

Day in the Life

OARC coordinates annual, and occasionally ad-hoc, Day in the Life ([DITL](#)) DNS traffic capture events. These involve many operators of significant DNS infrastructures—including root server operators, TLDs, and recursive operators—running packet captures of their traffic over the same 48 hour period. The data are uploaded to OARC where it is organized for use in research.

The DITL collections go back to 2009.

In recent years, the average DITL event collects around 12TB of compressed PCAP files from around 20 operators. We mostly receive data from root operators, with a small number of ccTLD operators and recursive server operators.

RSSAC 002 Statistics

The Root Server System Advisory Committee's publication [RSSAC 002](#) is the Advisory on Measurement of the Root Server System. It defines an initial set of statistics to be collected by root server operators from their systems. OARC collects the output of this reporting from each root server operator, daily, and maintains a history of these statistics available for analysis or review.

Zone File Repository

OARC maintains an historical [archive of zone files](#) which includes daily updates of the [root zone](#) going back to 1993, and weekly updates of several TLDs beginning at various times between 2009 and 2018.

Other Data

OARC also periodically accepts submissions of other data that may be relevant to researchers interested in the DNS:

- derivative data from research done on OARC’s other datasets
- data collected from OARC testing tools, such as Check My DNS
- DITL-like collections from outside regular DITL windows, such as occasional contributions from [AS112](#) server operators
- historical packet captures from OARC’s Open DNSSEC Validating Resolver ([ODVR](#)) which includes forwarded queries from the [DNS Privacy Testbed](#)
- Case Western Reserve University’s “Case Connection Zone” FTTH data
- other ad-hoc contributions of relevant data

System and Service Status

Routing Improvements

OARC was finally able to get a pair of new-to-us Cisco ASR9001 routers into production in our Fremont and Ottawa locations at the beginning of 2024, one router per location. This has several positive side-effects. First and foremost the random IPv6 route drops we were experiencing in Fremont have vanished, which improves the behaviour of numerous services.

The new routers also give us the capability to have multiple transit connections per site. We’re expecting to be able to turn up a second transit connection in Fremont this calendar year, contingent on available budget.

And finally, we’re now able to upgrade our upstream connections to 10GB, if any of our service providers end up requiring that. This isn’t something we’re expecting, but we now have the capability if any service provider decides to discontinue 1GB physical links.

DITL Update

The regular Day in the Life collection for 2024 took place in April, as usual. This year we received 13.5TB from 22 contributors.

New this year are recursive server contributions from Universidade Estadual Paulista ([UNESP](#)) in Brazil. They contributed from two separate recursive instances on different campuses, representing a total of around 7,000 clients.

As with the last few years, we are not running the “cleaning” process on uploaded PCAPs, which was intended to normalize the presentation of the data across operator contributions. The old process was having unintended side-effects, so we discontinued its use until we’re able to review it and solve its issues. Details are available in past reports and emails to the Members’ mailing list.

We’re planning to survey researchers and other potential users of the DITL data this coming winter about what facets of the cleaning process are still useful, and then troubleshoot and

redesign the cleaning tools as necessary.


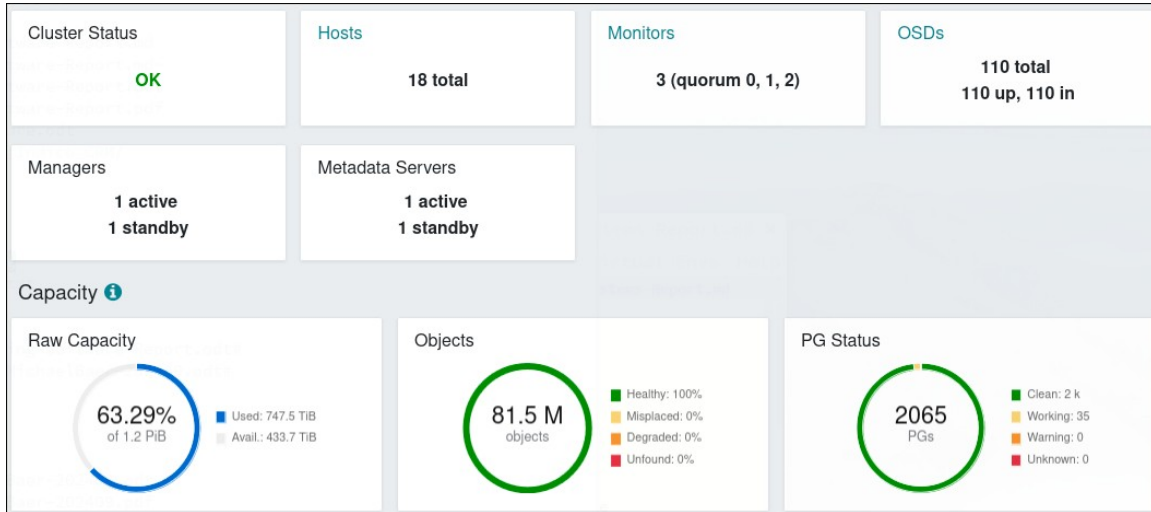


Figure 1: DITL: Contributors, Data Volume vs. Year

Ceph : Data Storage

With the additions of storage servers, Ceph became the data storage back end for all of OARC's available DITL data archive. It has redundant data storage and we plan to move the smaller data archives into Ceph moving forward. Figure 2 below is the current status of our Ceph storage cluster.



F2: Ceph Status: AGM 2024

Regular Operations

The vast majority of Operations work at OARC is spent on things like routine maintenance, OS and application updates, and support work for users of our services and systems.

Notable from among these subjects this year is that we have started a more serious move toward visualizing our systems. A number of “easy” moves have been made separate applications that have their own IP address for access (e.g. DNS servers) into containers. We’ll be following this up with shared address applications (such as web apps) in the coming year.

We’re using Incus as our container management application. It has several features we like, but does not directly support the sort of configuration discovery tools employed by application routers like Traefik. This means we’ll have to develop our own tools for automating configuration of an application router before we can move on to the web apps and similar applications.

OARC also added a Nextcloud server this year. The main purpose was to replace the previous method of shared file access OARC for staff, board, and occasionally external users. Nextcloud provides more control and flexibility for file sharing. It is also better supported on more platforms than the previous method. We also tried using Nextcloud’s shared calendar system. Unfortunately, it had some compatibility issues (e.g. with Google) and didn’t work well with most peoples preferred calendaring usage. We moved back to the previous OARC Google calendar.

Portal

The support for OARC’s portal software moved from Jerry Lundström to Michael Baer this year. It is the location where OARC Member’s can manage their membership and find other member’s

contact information. It is fairly stable at this point and with the additional new work that Jerry is doing, particularly with [Crunchy DITL](#), we moved it to operational support.

There were several minor revisions of Portal released since the last AGM. They included:

- Numerous text changes: clearer text, copyright update, etc.
- Mail handling improvement: Attenuate auto-replies for bulk mail.
- Changes/Additions to various exports to better support staff.
- Bug fixes

AGM 2024 - Software Report

Jerry Lundström

Project: Crunchy DITL

This project is funded by Verisign and aims to try and make DITL data easier to research on by making it available in a Big Data Tool/Data Lake.

Last year we evaluated different platforms and decided to go forward with ClickHouse. So early this year we acquired the hardware needed, installed it in Fremont (our main site) and began designing the import software.

DITL 2020 was chosen, some 16TB+ compressed PCAPs, and we picked dnsjit as the tool to parse them. For that, there's a couple of new modules, like zmpcap - memory mapped compressed PCAP reader and a ClickHouse client using ClickHouse's own C++ library. The dnsjit scripts were then controlled by a bit of python and ZeroMQ along with Flask/socket.io and some HTML/JavaScript to monitor progress.

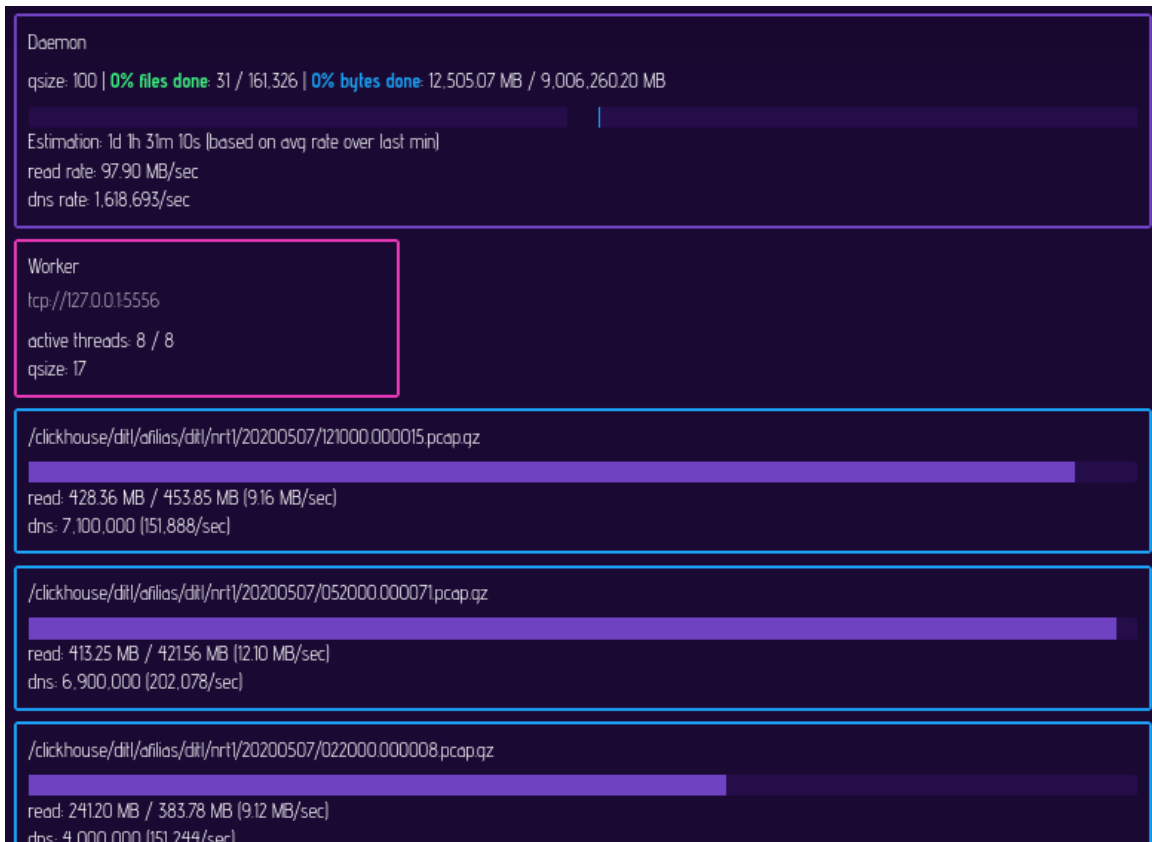


Figure 3: Crunchy DITL

Once up and running it could parse all the compressed PCAPs, split across two servers, in about 6 hours, but just parsing. Next came creating the database schema for ClickHouse and we went with a few variants, DNS header bits as a uint16 or group of booleans, QNAME reversed etc. Each variant import took around 5 days.

Stats: - DITL 2020 RAW compressed PCAPs, 16.3 TB - 297.93 B (297,925,409,197) DNS records imported - ~11.5 TB compressed, ~50.7 TB uncompressed in ClickHouse (per schema variant)

Next up was for a few researchers to try it out and beside people from Verisign we decided to invite a few more. The conclusion of this evaluation is still pending at the time of writing this report but hopefully we can share more during OARC 43. And during that conference I (Jerry) hope to be able to demonstrate it live in the hallways during breaks.

Open Source Releases

Here's a brief description of all releases since last AGM, all our open source software can be found on GitHub: [DNS-OARC's Github](#)

dsc-datatool, two releases, latest is v1.4.2

Fixed issue with InfluxDB quoting, was missing to quote the quote character. And fixed issue with IANA's IPv6 parameters file, dsc-datatool expected a RIR in the Designation field but IANA recently added a title for SRv6 reservation which caused an exception.

dnsjit, two releases, latest is v1.4.0

Added `core.object.dns:reset()` to reset DNS objects so they can be reused correctly. Example scripts have been updated to reflect this along with correctly setting `includes_dnslen` for TCP packets.

New example converter `pcap2tcpdns` has been added, it can be used to create binary files for replaying with `dnsperf -B`.

Added the new module `input.zmmpcap` to read compressed PCAPs using `mmap()`.

Removed blocking of signals during start up.

In `core.thread`, the `push()/pop()` now uses `double` instead of `int64_t` to match the conversion table of LuaJIT and to support floating point numbers.

The `dumpdns` example has been updated to support reading PCAPs using `mmpcap` or `zmmpcap` and reading PCAPs from `stdin`.

In `input.zpcap` the support for reading `zlib/lzma` compressed PCAPs has been added along with fixing a bug where it incorrect checked if it had read enough.

Fix a potential memory leak in `lib.trie`.

dumumd, two releases, latest is v1.4.0

New stuff: - Add `dumdohd`, reflect DNS-over-HTTPS using `nghttp2` example server - Add `-Q <num>` to set listen backlog - Dockerfile - Add `clang-format` and `format` code - Add `-o flip-qr-bit` to flip the QR bit when reflecting - Add `-A/-R` (reuse `addr/port`) to `dumdohd`

Bugfixes and tweaks: - Update README about compiling environment - Fix #20: Flush stdout on stats - `dumdohd`: Remove deprecated OpenSSL functions - `dumumd`: Ignore SIGPIPE, can happen during `BIO_write()`

dnsperf, one release, v2.14.0

In “Connection Statistics”, `reconnections` has been renamed to `connection attempts` and now includes both the initial connection attempt and following reconnections.

If supported, threads will now be named after what they do, such as “`perf-send-`” and “`perf-recv-`”.

The TSIG context was shared between all threads and would cause a crash if more than one sending thread was used. This has been fixed and TSIG contexts are now per thread.

dsc, two releases, latest is v2.15.2

Fixed client subnet indexer which overwrote the mask options during initialization so the configure option `client_v4_mask` and `client_v6_mask` was never used.

This releases fixes 3 issues detected by code analysis tools:

- File not closed and memory not freed during error while loading known TLD file
- Label buffer should be static
- Unsigned difference expression

Also updated the builtin known TLDs based on PSL.

golang-dns-server-doq, one release, v0.4.0

Updated to Go 1.22 and fixed handling of read/write deadlines, writing was timed out immediately.

PacketQ, two releases, latest is v1.7.3

Fixed a memory alignment issue, the handling of TCP segments and various other issues reported by CI/code analysis tools.

dnscap, one release, v2.2.1

Fixed the BPF to correct capture ICMPv6, a typo in the filter had it matching “ip” and not “ip6”.

Software Packages

There’s been a bit of change to Debian packages since last AGM. We no longer build these packages ourselves, instead we use SUSE Open Build System (OBS).

This allows us to build for most versions of Debian’s distribution, such as oldoldstable and unstable etc.

For full details about our packages, please see [DNS-OARC Packages](#).

End of mailing-lists?

We have decided to shutdown the mailing-lists for DSC and dnscap due to community inactivity, we have basically just been posting release announcements for the last couple of years.

They have both been closed down, as in you can’t post or subscribe to them anymore, but the archive will still be around.

So where can you talk about OARC software then? Well, in the OARC Software channel on our Mattermost of course :D come join us!

[DNS-OARC Mattermost: Community: oarc-software](#)