

# Engineering Report

DNS-OARC AGM 2023

Michael Baer – Systems Engineering  
Jerry Lundström – Software Engineering  
Matthew Pounsett – Systems Engineering

This year we are combining reports from the various corners of Engineering into a single document. We expect this to help with minimizing duplication, and keeping the report more concise.

As always, we will primarily be reporting on the major changes and events over the past year. For a deeper dives, or more detail, please see our other reports released throughout the year as well as Jerry's Development Updates, and release notes from our software. And, as always, we are happy to answer questions or requests for more detail.

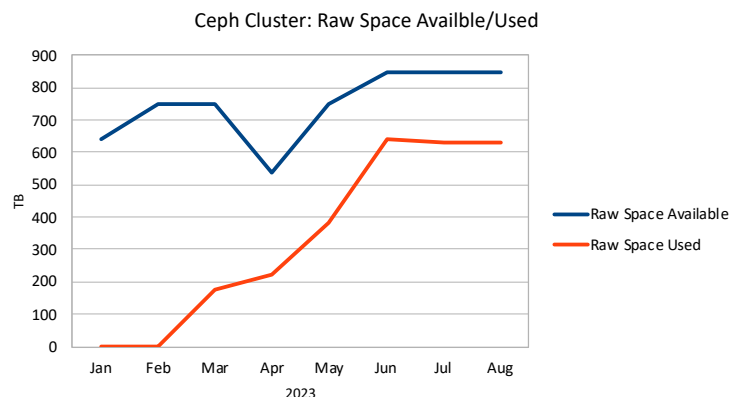
## Systems and Network


### Ceph 2023

Since the last AGM, OARC's new [Ceph](#) data storage infrastructure has become a working cluster and we've completed the majority of the data-copy from our aging NFS storage system.

In March of this year we released our [Ceph Deployment Plan](#), which was the result of much discussion about the order of operations and validation we would be using to safely migrate all of our data to the new cluster. While we found in practice that some steps were best done in a slightly different order, or in parallel, we have kept to the plan quite closely. The remaining steps from that plan begin with the recovery of offline hardware, which we expect to complete over the next few months. We will also be adding two new storage servers before the end of the year, to account for required capacity growth since we initially purchased the hardware.

As we move data over to the new cluster from older NFS servers, we are beginning to remove access to the old copies of the data. Researchers and other users of the analysis systems should pay attention to the





Message of the Day shown at login, in order to be aware of datasets that should be accessed by a new filesystem path. Next up on the list of old paths to be made unavailable is Pool9, as that server is part of the new cluster hardware "borrowed" for the NFS infrastructure during the transition. It is due to be repatriated to the cluster in the next couple of weeks.

With the exception of the unknowns of data recovery on NFS pool5, which our previous reports have noted is offline due to hardware issues, OARC is on track to have all of the Fremont NFS data store transferred and accessible for analysis from the Ceph cluster by the end of the year.

## Configuration Management


Late last year we finally began to roll out active configuration management on OARC's systems. [Saltstack](#) joins [FAI](#) (Fully Automated Installation) in our efforts to move toward having easily replaceable systems.

OARC's systems have historically had an unusual amount of diversity in the way they are configured. The lack of consistency has been an issue we have gradually been correcting over the past few years, and has reached a point where automating the configuration is now feasible within our limited time budget. Salt is already responsible for the configuration of nearly all of our base OS configuration and package installation, for the configuration of many applications, and is used as an orchestration tool for routine tasks like system patching and checking hardware inventory.

Although this will be a very long road, the eventual goal is to reach a point where failed systems can be re-imaged and replaced within a few minutes of having replacement hardware on hand. We would also like to see all internally developed applications installed, updated, and configured as part of a CI/CD pipeline.

## Routing Improvements

We reported last AGM on the purchase of a pair of new routers to replace the extremely old hardware in place in Fremont and Ottawa. At that time we were searching for a contractor to help with the conversion and cleanup of the configuration, as we did not have time in our operations calendar to do that work with existing staff. It took many months, but we did eventually get a contractor in to help with that deployment, and got the project moving properly this summer. We are finally on the verge of rolling them both into production!



# DNS-OARC

Domain Name System Operations Analysis and Research Center

Some other delays and priority conflicts have continued to set the project back, but we now expect that we should be moving both routers into production in mid November.

With the new router in Fremont, we will finally be in a position to add redundant transit, and to join a second large IXP ([FCIX](#)) at the site. Both expansions have been on our wish list for several years, and should be very easy to deploy once we have hardware in place that will support the additional interfaces and BGP sessions. Our plan is to deploy both in the first half of 2024.



## New Mastodon Instance


Earlier this year, OARC joined in support of the Fediverse by launching the [MastoDNS.net](#) Mastodon instance. Use of the platform continues to grow, and in addition to many humans we are now also host to the [@diffroot](#) bot, with one other in the wings hopefully arriving soon.

## AS112 Project Advancements

We have recently released a complete redevelopment of the AS112 Project [web site](#), which was generously funded by the [American Registry for Internet Numbers](#) (ARIN) as part of their [Community Grant Programme](#).

As part of this redevelopment, much of the supporting information about the AS112 Project has been rewritten, and the site has been reorganized to make that information easier to find. We have also modernized the HTML to include things like proper metadata for search engines, reactive design for viewing the site on different devices, and accessibility for assistive devices like the screen readers used by visually impaired users.

The operator listing has been converted from the rather messy hand-edited HTML table to structured data (JSON), and now properly supports operators with multiple instances—a setup that was inconsistently hacked into the table in varying ways, before. We also now download and cache information from [PeeringDB](#) on a daily basis, covering AS112 nodes that peer directly at IXPs, and the site supports cross-linking between Operator and IXP lists. The new structured data for the Operator listing makes that data more easily machine-readable, which can also support future enhancements. For example, regular automated surveys of the unicast addresses and site statistics could be done to look for issues and/or old listings that should be removed.



The new site is also built with translation in mind. Although we do not have the resources at this time to fund any translation work, the entire site was built on gettext-style translation tables (using Python [Babel](#)) which will make such translations easy to deploy if we are able to put aside those resources, or if volunteers come forward help with that work.

Look for a blog post in the next few weeks going into more detail about the development work.

## DITL Update


The Day in the Life data collection for 2023 took place as usual in April. This year we received around 12TB of packet captures from 18 contributors, which is about average for recent years.

We are not running the "cleaning" process for the 2023 DITL. As we have noted in a number of different reports and other communications with members over the last few years, we have been seeing an increasing amount of unexplained data being discarded during this process. Please see previous reports for more detail about these issues.

We theorize that the problems stem from modern uses of the protocol not accounted for in the tools used to perform this data cleanup, but we have lacked the time and specific documentation about those tools' operation to confirm this hypothesis, or come up with a more satisfactory explanation.

During the upcoming year, we will be reviewing the cleanup process in more detail. We will be looking at each of the known, expected outcomes of the process and identifying the parts of the existing tool chain that accomplish the outcomes that we know are useful to researchers. The goal will be to discard parts of the toolchain that are not making explainable changes, or that don't meet the more narrow set of goals, and to more clearly document—or possibly rewrite and improve on—the parts of the toolchain that are. We hope that the outcome will be a more useful to researchers than either the unprocessed "RAW" uploads, or the last few years of "CLEAN" data.

We are still considering whether it makes sense to re-process every previous year of DITL data for consistency with future datasets, or to leave them as-is for consistency with previous explorations of those data. Input from researchers on this would be greatly appreciated, and will be specifically sought out once we reach the point when a decision must be made.



We also have redevelopment of a lot of our server-side management of the actual DITL events on our roadmap. This will involve building internal tools to update how we manage credentials for contributors, report on DITL collections in-progress (for contributors to verify that their uploaded data is being received), and simplify and improve the metadata we publish about DITL collections to researchers after a collection event is over.

## Regular Operations


Although little of this work rises to the level of public interest necessary to include details in many of our reports, the vast majority of our time is spent on things like routine maintenance, application and OS updates, and support requests. In addition to migrations between major OS releases, there have been significant updates to both internal and public-facing applications, and whole new internal tools deployed to assist staff, such as our new self-hosted, shared corporate password store using [Bitwarden](#).

We have also undertaken a significant amount of work in the datacenters where we have equipment, to catch up from several years of limited travel and reliance on remote hands. There has been a significant push to update inventory, documentation, and labeling, and we have cleaned out a huge volume of old hardware awaiting recycling and/or secure destruction.

OARC has also implemented a new hardware replacement budget. In the past we have mostly replaced hardware only upon failure, but as of the 2023 budget we are now attempting to do routine replacement in a way that more closely matches industry standard. Although it will take several years for all of OARC's hardware to see replacement under this policy, and while we expect there to be some slippage due to typical not-for-profit budget constraints, we will be attempting to meet this schedule:

- servers to be replaced every 5 years with new hardware
- routers and switches to be replaced every 8 years with refurbished hardware
- other support infrastructure (console servers, power distribution, etc.) replaced every 8 years with new hardware
- warranties extended beyond the default only for critical hardware that is difficult to replace quickly (e.g. database servers)

Over the next year we plan on deploying our own [NextCloud](#) instance to simplify a lot of our internal file sharing, and possibly move some other cloud services (like calendaring) in-house. We are also considering the possibility of relocating our guest-hosted Ottawa site to a location we can control more directly.



## Software Engineering

### New DSC Grafana example dashboards!


In the first quarter of 2023 we finalized new example dashboards for [DSC data in Grafana](#)! Have a look at the links below which show [DSC](#) data coming from [Check My DNS](#).



The “[Statistics](#)” dashboard contains some simple stats on packets, queries, responses and response codes (RCODE), and their relation as a difference graph such as “queries vs responses”.

Then there’s a “Query Failure Rate” graph which shows a percentage of responses seen with an error related to total queries seen. Note, there is no relationship between query and response here, but it can still show some interesting stuff! What we’ve seen is that sometimes it goes way beyond 100%, and it looks like someone is sending bursts of faked/spoofed failed responses. Possibly some cache poisoning attempts or something else.

Last on the dashboard we have two graphs showing queries with DO and RD bit set and their ratio of all queries seen.



# DNS-OARC

Domain Name System Operations Analysis and Research Center

The second dashboard, “[Locations and Queries](#),” uses the new Geomap plugin to display where queries are coming from based on country code lookups using MaxMind databases along with other stats based on the query.


The “RCODE per IP/Net (top20)” graph is something we're very proud of! This was one of the graphs we couldn't recreate a few years ago when first using Grafana for DSC data. This is because you need to turn the data on its side, and this is now possible thanks to Grafana's new transformations. Similar things can be done with InfluxDB's new [Flux query language](#), but we haven't had time to get into that yet.

We hope these example graphs are helpful for people! And if you want to share your own graphs/dashboards we've created a repository, [dsc-datatool-grafana](#), for that.

## DNS Hackathon 2023 — Connect to port 53!

During May 20–21 (the weekend before #RIPE86), 42+ hackers gathered from all corners of the world to work on DNS with the challenges of; Sustainability, Security (& Privacy) and Sovereignty.





# DNS-OARC

Domain Name System Operations Analysis and Research Center

We had 6 teams and 6+ projects:

- **DApper**: DNS and the Application layer
- **diggin-in**: experiments in meta-language creation for measuring DNS
- **DNS-Today**: DNS Traffic Monitoring and Classification
- **DNSSEC-Bootstrapping**
- **DNS-oops**: work on the personal draft for DNS out of protocol signaling
- **DNS-Home-Appliance**: run a low-cost but fast DNS home appliance
- **SustainabiliTeam**

The hackathon was hosted by us (DNS-OARC), Netnod and RIPE NCC with the grateful support from our sponsors SIDN, NLnet Labs, ISC, Gcore and ICANN.



## Awards

We had three categories of awards, in no particular order:

- DNS-oops was awarded “Most useful project for operators” for their work on out of band signaling, in practice, triggering BGP! This project was also be presented at DNS-WG RIPE86, check [slides](#) and [video recordings](#).
- DIG-alicious was awarded “Best teamwork, exemplified hackathon spirit” for their work on RIPE Atlas data compression and pre-scripting!
- DApper was awarded for “Team tackled a difficult/important challenge and progressed a potential solution” for their work on the DNS and the application layer!

For more information about #DNSHackathon2023, project results and presentations, see [DNS Hackathon 2023](#) organization on GitHub.

## Portal — Event Discounts

Hopefully you've signed up for the AGM using the new Events system we added to the Member Portal.



OARC Portal

[Links & News](#)

[Events](#)


[Services](#)

[Organization](#)

[MyAccount](#)

This is a huge step for us to automate and have a simple way to manage discounts to our conferences for our members. Hopefully you also found it simple to use! If you have any issues, questions or found any bugs with this new feature please don't hesitate to reach out to us!





## Development Updates

Here is a list of updates since AGM 2022:

- [Development Update #2211](#)
- [Development Update #2302](#)
- [Development Update #2305](#)

And if you happen to have a furry master at home that you want to be featured in upcoming development updates then please send Jerry pics.

Software releases:

- 22 releases of open source projects
- 4 Portal releases
- Over 220 issues and pull requests

For release notes please see <https://github.com/DNS-OARC>

## Upcoming!

An exciting new project related to DITL data has just started!

We will try and build a prototype for processing DITL data using modern and extendible big data systems, together with a structured query language. Gone will be the days when everyone needed to code their own analysis software for the PCAPs, hopefully! 😊

This new system will hopefully make DITL analysis much more convenient and available for any member that wants to poke at the data! And you will be able to get a reference to the source for your results if you want to take a closer look at the raw material.

The prototype will include a limited set of DITL and we hope to have it ready for testing within a year. If this project is successful, then we will seek more funding to add capacity and processing power for the whole DITL data set that is currently 175 TB and grows around 12 TB per year!