Systems Engineering Update
OARC 38, Philadelphia

Michael Baer
Matthew Pounsett

2022/07/20

Contents

1 Introduction

The Systems Engineering focus for the past half year has continued to be on clearing the backlog of issues and updates, in order to clear time for more forward-looking changes and improvements. The addition of Michael Baer to the team this Spring is already having positive results. And we've also been able to return to something resembling normal travel which allows for on-site visits and less reliance on remote hands work.

We have been redistributing the backlog between the two of us, which means that the list of projects seems much more manageable now. Activity will soon slow for the summer, but we're anticipating more externally visible improvements in the coming months and into the fall.

2 OARC Services Overview

2.1 Data Archiving

OARC maintains a large store of multiple data sets.

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.

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

2.2 File Servers and Storage

OARC's datasets are stored on several discrete file servers. Five of the file servers, located in Fremont, California, have 343TB used of their 505TB of capacity. Two of these have multiple filesystems, marked as A and B in the chart below. The sixth file server, located in Ottawa, Ontario, is an off-site copy of a selection of datasets from the first five servers.

We are in the process of building a new Ceph storage cluster, discussed in detail in previous reports, and covered again below in section 3.5, which will replace the old discrete NFS servers. Due to their age, the old servers have been having increasing trouble with stability in the last few years. We are currently "borrowing" one of the new storage servers as a temporary addition to the older servers in order to cover for lack of available storage on the couple of stable servers we still have.

There is one file server, not listed below, which lost its filesystem due to multiple disk failures in too quick succession for us to recover the RAID. Another, FS5, listed as temporarily offline, experienced what we believe to be either a CPU or memory failure which is preventing it from booting at the moment. We will cannibalize other servers for parts in order to boot it when it comes time to migrate its data to the new Ceph cluster.

Server/Volume	Used	Capacity	Notes
FS2a	36TB	42TB	
FS2b	75TB	125TB	
FS3	34TB	42TB	
FS4	72TB	84TB	
FS5a	69TB	84TB	Temporarily offline
FS5b	33TB	42TB	Temporarily offline
Stor09	46TB	86TB	borrowed from new cluster
FS6	117TB	121TB	Located in Ottawa, Canada

Each file server uses either ZFS (RaidZ2) or XFS over software RAID for its filesystem to provide redundancy within the file server. Each dataset is intended to be stored on more than one file server in order to create cross-chassis redundancy of data; and, due to the server in Ottawa, some datasets currently have copies on three systems.

Due to the loss of the `FS1` server, not all datasets have their expected number of copies available, and with `FS5` offline, some datasets have either no copy accessible or are only accessible on the Ottawa server.

All capacity numbers above are the filesystem capacity, rather than the raw size of the disks in service.

2.3 Data Analysis Servers

OARC maintains four Linux shell servers with access to the above data sets. Three in Fremont, CA (`an1`, `an2`, `an4`) and one in Ottawa, ON (`an3`). Members and Supporters who have signed a Data Sharing Agreement and request access are given accounts on these analysis servers, which they can use to do research into the DNS using any of OARC's datasets.

> **Note Well**: No data, even derived data, may leave OARC analysis systems without express written authorization, in compliance with the Data Sharing Agreement. Contact admin@dns-oarc.net first, *always*.

3 System and Service Status

3.1 General Condition

OARC's network and systems have been mostly stable since the last Systems Engineering Report. As before, there remain a number of systems and services that should be changed/upgraded to require less care and attention. While we still need to clean up and fix a number of high-maintenance systems, the list is shrinking. And there has not been any major disruptive system surprises.

The pandemic disrupted much of our regular hands-on work. But we seem to be on track to getting back to normal with respect to on-site visits and physical maintenance.

3.2 Ottawa Site Relocation

For a number of years now, CIRA has been generously hosting some of our "off-site" systems in one of their POPs in Ottawa, Canada. They informed us this spring that they would be relocating that site, and that we would need to move along with them. The site move was completed with only minor disruptions to the services there, and in the process we managed to improve several things with the site that had been annoyances of one degree or another.

The local network is now using the same internal standards for configuration as our main site, segregating off different types of internal traffic. Renumbering of the private address space is now also complete, removing the remaining overlaps with private address space used in other sites. This will allow us to tunnel our private VLANs between sites, which will improve access to things like virtual consoles.

We also added our own transit connection to the site, so we are no longer piggy-backing off of CIRA's transit connection. The Ottawa router needs to be upgraded, but once that is complete we will be moving from static routing to BGP, and will be seeking an IPv6 assignment from ARIN. For IPv4 addressing we are considering a few options, including the possibilities of splitting our existing /23 into separate networks, and speaking to an address broker about an additional /24.

A new router for Ottawa has been obtained and will be shipped to Ottawa at our next Fremont visit, which is expected sometime in August. There are some questions of import process which need to be answered before that can be executed. We hope to get that installed during the summer, and will be investigating the possibility of secondary transit and new address blocks in the fall.

A second storage array has been added to the site, which is acting as a backup for `fs6`, our off-site copy of some of our older data sets. The future of our Ottawa copies of analysis data is still under consideration, but in the short term this alleviates concern we had over some datasets which no longer had a second copy in our network, due to the failures of `fs1` and `fs5` in Fremont.

3.3 Networks, Routing, and Routers

We finally have an end in sight to the long-running issue with random IPv6 routes not being installed in our Fremont router. There is new equipment! We have a router ready and waiting in Fremont, which will be configured this summer, and should take over production routing before the end of Q3, possibly even before the end of September.

We expected it to already be in standby for production service, however the timing of a shipment of optics didn't work out, and they didn't arrive in time for our last on-site visit. Our next opportunity to install those optics comes up in the first half of August.

Once the new router is in place we expect the v6 routing issue to go away. We will also finally be in a position where we can carry the routes from a second transit connection. We already have some plans in the works and will be ready to jump on that as soon as hardware permits.

3.4 Updates, Upgrades, and Upshot

Several OARC system software systems have been upgraded since the last Systems Engineering Report:

Drupal: Drupal is the software that maintains OARC's www.dns-oarc.net website. It was updated to the most recent supported v9.4 version.

Indico: Indico is the software OARC uses for conference and event handling. It was also updated to the newest supported v3.1 version.

3.5 File Servers

The Old Stuff

Not much has changed with the old file servers. They have been relatively stable since the last Systems Engineering Report. But they are still aged to the point of replacement with the associated higher risk of failures.

The New Stuff

The new Ceph distributed file system that is to replace OARC's current NFS data storage system is still in process. Currently, the hardware is racked with the management and access network provisioned. The software management to allow auto creation of the nodes is being debugged/refined. The next steps include: Ceph management, monitoring and storage node creation; data network provisioning; and data transfer.

3.6 Day in the Life Dataset

The 2022 Day in the Life collection was completed in April. We received 11TB from 20 contributors, which is down 2TB from the previous year. The raw data (as uploaded by the contributors) was released to researchers May 19th, in keeping with our new procedure for handling DITL data releases.

Our hunt for volunteers to assist with evaluating and revamping our post-processing of the DITL data did not yield any results. As such we've had to make do with the limited internal resources for trying to identify and correct any issues with that processing. As a result, the first pass on cleaning up those data wasn't complete until recently, and unfortunately with workshop prep taking precedence, there hasn't been time to review the results. It's expected that review will take place the first week of August, and it seems likely that at least a partial release of cleaned-up data will take place in the second week.

We're hoping to continue to improve things over the coming year, and hopefully the data cleaning process will be smoother for the 2023 collection.