OARC Systems Update

OARC 29, Amsterdam, October 2018

Contents

1 Introduction

Since OARC's Systems Engineering role went through a change of hands without the opportunity for proper handoff, it's likely this Systems Engineering Update will not maintain the usual continuity with previous updates. An effort has been made, however, to bring forward and address any open issues from previous updates. Questions about this update should be addressed to admin@dns-oarc.net.

This written report is an experiment in finding a new, more detailed format in which to report status, changes, and plans to members. It is intended as a companion to the oral report given at the OARC Workshop. Feedback on its usefulness is appreciated.

2 Data Archiving and Analysis

OARC maintains a large set of data comprised of:

- annual and occasional Day in the Life (DITL) DNS packet captures primarily from root server operators, TLDs, and recursive operators
- DNS Statistics Collector (DSC) data submitted by members
- RSSAC 002 statistics
- a Zone File Repository including a Root Zone Archive going back to 1993
- other data submitted by members and other participants in OARC

2.1 File Servers and Storage

OARC's datasets are stored on six file servers. The first five file servers, located in Fremont, California, have 358.82TB used of their 536.27TB 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 the first.

In September of this year a new 125TB filesystem was added to Fs2, comprised of 20 12TB disks, bringing OARC's data capacity in Fremont to over a half petabyte, and global capacity to 0.64PB.

File Server	Used	Capacity
Fs1	118TB	121TB
Fs2a	36TB	42TB
Fs2b	0TB	125TB
Fs3	34TB	42TB
Fs4	72TB	84TB
Fs5a	69TB	84TB
Fs5b	33TB	42TB
Fs6	117TB	121TB

Each file server uses either ZFS (RaidZ2) or XFS over hardware RAID to for its filesystem to provide redundancy within the file server. Each dataset is stored on more than one file server in order to create cross-chassis redundancy of data. This means that the total size of all unique datasets is roughly half of the 475TB indicated above.

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

2.2 Data Analysis Servers

OARC maintains four UNIX shell servers with access to the above data sets. Three in Fremont, CA (an1, an2, an4) and one in Ottawa, ON (an3). Members and participants who have signed a Data Sharing Agreement and request access are given accounts on these analysis servers.

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

2.3 Root KSK Roll DITL

OARC has been running an extra DITL event this year surrounding the root KSK roll. Since the KSK roll has just completed, the data upload is still ongoing. As of writing, OARC has acquired 5.5 terabytes of new DITL data from 12 contributors (10 root servers, 2 TLD operators).

Because of the way some contributors manage their capture data, much more data from more contributors is expected to be uploaded after the end of the capture period.

2.4 New Dataset

In the next week or two, OARC is expected begin receiving a live feed of authoritative DNS server query logs courtesy of John T. Kristoff and the DataPlane project. The dataset is being provided as part of DePaul University's supporter status. This is the inaugural dataset of a new Syslog-over-TLS submission transport OARC is launching.

The syslog data will be available from Fs3's collection.

3 System & Service Status

3.1 General Condition

Most of OARC's services were originally deployed in a volunteer or early startup environment. This meant that minimising development time and other deployment efforts were the high priorities, with predictable effects on long term maintainability and documentation. This was the correct approach in OARC's early years when systems engineering, software development, research, and administration fell to a group of volunteers or to a single individual employee.

However, there is little evidence that any effort has been expended in recent years to bring OARC's systems and services up to date with current best practices, or to update service architecture to reduce maintenance effort. Documentation has been improved, but still mostly serves as a reminder to those already familiar with the deployment of OARC's systems and services.

As the number of systems and services has grown, the maintenance effort required has also increased dramatically. Most services do not have sufficient documentation for an unfamiliar operator to troubleshoot, are not properly monitored, and continue to use outdated techniques or technology.

Remediation of these issues is discussed in the Future Work section, below.

3.2 OS Distribution Choice and Configuration

Earlier this year, steps were taken which increase the engineering workload by replacing all Debian systems with Devuan, a systemd-free fork of Debian Linux. On the surface this seems like a benign change, however it is not. Since most major Linux distributions have adopted systemd, third-party packages typically require it, making them non-installable on Devuan systems. The Devuan distribution is still quite young, and as a result its package repository is still quite small. This results in a very large number of applications which must have custom packages made in order to install them on OARC systems, which is not scalable for a small shop.

The OARC systems generally overuse an `/etc/rc.local` script for service startup and system configuration, where it is normally intended only as a last resort or for temporary changes. For example, the IPv6 configuration on systems built before June of this year is all in `rc.local` rather than in the OS distribution's standard network config files. This is also a place where the custom-compiled applications are a problem, as they are all started out of `rc.local` rather than using the distribution's standard INIT system. The result being that it is more difficult than normal to start and stop applications outside of rebooting the system. The *average* length of the `rc.local` file on OARC systems is over 85 lines, with the largest being 138 lines.

Remediation of these issues is discussed in the Future Work section, below.

3.3 Monitoring

OARC uses Nagios to monitor the availability of its systems and services. As of June, the Nagios web UI had never been configured. When the web UI was fixed, over 40 active alarms were found. Most of those alarms were for monitors that had never worked as intended, or had never worked at all. Today that number is down to seven; the remaining alarms are primarily for near-full storage for static data sets (e.g. on the large file servers), as well as a couple of known issues mentioned below.

The configuration uses a number of home-grown monitors that have open source counterparts (mostly available as part of the default install) which should be replaced. It also inconsistently uses templating, and host and service groups, making reconfiguration complicated. There remain a large number of services that are not monitored correctly, or not monitored at all.

At a minimum, the Nagios configuration is going to be reimplemented from scratch, using standard tools as much as possible, and more predictable configuration patterns. More modern monitoring tools are also going to be evaluated for their suitability to OARC's needs and budget.

3.4 Backups

The backup software in use for OARC's systems is currently far out of date. The version being used was released in April 2013; the most recent version in the same major release is nearly two years old, and a new major revision has been available for the last 18 months. At this point, significant testing will be necessary before putting an update into production, to ensure that the configuration is still valid.

As of June, the Backup system's web UI had never been secured behind ACLs or even HTTP Basic authentication, and instead access was controlled by shutting down the web server when not in use. This has been corrected.

The file store on OARC's backup system has been running at 98% at least since the beginning of the year. In recent months, as services are cleaned up, and systems which were not being backed up have been added, storage has begun to fill and backups to fail on occasion. In order to alleviate the problem, the age of backups being kept has been reduced to two weeks while we evaluate options for a long term solution.

3.5 DANE Test Pages

The DANE test pages were deployed in 2016 to demonstrate different types of behaviour that could be experienced when browsing DANE-enabled web sites. The web server which hosts the DANE test pages has an issue with its HTTPS module which makes it incapable of negotiating SSL/TLS. The web server in question appears to be a custom compile of Apache, for unknown reasons.

OARC's monitoring has been alarming about this issue since mid-December 2017.

3.6 Zone File Repository

Aside from the major failure encountered in July, which was extended by a lack of documentation (discussed below), we continue to have problems with the AERO and ASIA downloads, which have been failing since May 2016. This appears to coincide with the renumbering of systems as they were moved from ISC to Hurricane Electric, and it is suspected the problem is an ACL on the FTP servers which provide the zone download service.

Several emails have been exchanged with Afilias in an effort to correct this, but we have not yet been put in touch with the correct party to troubleshoot the problem.

3.7 DITL Collection Software

The DITL collection software suffers from some of the same problems as ZFR. Written at a time when the priority was to get something working and out the door, it has minimal documentation, and lots of moving parts and manual steps involved in setting up, managing, reporting on, and finalizing a data collection event. There have also been several constructive criticisms received in the last month regarding the state of the client side software.

At some point, re-engineering of the DITL collection client and server software will be necessary. This may not require any changes to the client/server interface, which will be good news to operators who have written their own collection software.

3.8 Mailing Lists

An issue was raised this summer about OARC's mailing list configuration not being friendly to some more restrictive configurations of DKIM and DMARC. There are some easy fixes for DKIM which we have not yet had time to implement, partly because they involve list behaviour changes that should be announced well in advance to the list memberships.

Changes to make Mailman lists more friendly to DMARC seem, at first glance, to be more difficult but we have not yet had time to evaluate how difficult they actually are, or whether to pursue them.

3.9 Physical Operations

OARC's cabinets at Hurricane Electric are reaching capacity due to the way that the physical network and systems are laid out. We currently have three switches, covering the public and private networks, installed in two of the five cabinets at Fremont, with inconsistent use of VLANs. With this layout, all network cabling must be fed between cabinets. We are also, in some cases, running power cables between cabinets because insufficient rack space was left for managed power bars. This has choked the exit points for most cabinets to the point that it is extremely difficult to run cables for new systems.

A budget and work plan is being developed to re-deploy the existing systems into the same cabinets in a way that takes future expansion into account, putting a switch in each cabinet so that only uplink cabling and fibre needs to pass between cabinets. It is hoped that the plan will be ready by the end of the year, with execution taking place as early as possible in 2019.

3.10 Analysis Server Resource Usage

During this summer, OARC has experienced several periods where analysis servers were temporarily unavailable due to resource overuse by users. The most common issue is storage in the home directories, which has been hovering near full on all servers at least since the spring. Several emails have been sent to users who have very large amounts of data in their home directories which appear to be no longer in use, and some have been kind enough to delete old data to alleviate that problem.

We have also had issues with all physical and virtual memory being consumed, and all CPU being consumed, which prevents administrative logins and interferes with monitoring and measurement of the hosts.

Later this fall or early winter we will be implementing simple quotas on storage, memory, and CPU in order to prevent complete exhaustion of these resources by non-administrative users. These will be implemented as umbrella quotas, covering all concurrent usage by all non-administrative accounts, so as to avoid unnecessarily limiting any individual user should they be the only one currently using a system.

OARC is considering the possility of implementing more directed quotas in order to promote fair usage of resources between analysis users, but no detailed proposal has yet been considered. We will be monitoring the situation over the coming months to see if proceeding in this direction is necessary.

The possibility of deploying an "exclusive" analysis server, suggested at OARC 28, is still under consideration, but is temporarily on hold due to budget considerations, as is the plan to expand local storage on the analysis servers by attaching JBOD drive shelves.

3.11 Election Software

OARC uses an open source application called OpenSTV to run the board elections. This package has been unmaintained since 2011 when it was taken commercial, and the web site for it has recently disappeared as the commercial version has been withdrawn and replaced with a SaaS site. The latest open source release of OpenSTV is archived by a third party on GitHub.

OpenSTV depends on a windowing library called wxPython, and uses function calls which are deprecated in current versions of the library. The most recent version of wxPython that OpenSTV can use in turn depends on deprecated c++ library calls, making it difficult to compile on current operating systems.

The election this year will be run on a Ubuntu 14.04 VM, which appears to be the most recent OS capable of running the software using easily installed packages.

It is considered important that OARC be able to run election counting software itself, rather than relying on a cloud-based application. Unfortunately, it appears that there are no more currently maintained, open source voting applications available for download.

3.12 Addressing Overlap

Currently, the "back end" networks in the Fremont and Ottawa data centres are using the same RFC1918 /24. We intend to set up a GRE tunnel between sites to allow more direct communication, but this overlap in addressing complicates matters.

The Ottawa site will be renumbered, and as the systems there are not public facing it is not expected that anyone will notice any downtime or service interruptions.

3.13 Other Issues

There remain two dozen other alerts, errors, misconfigurations and other apparent problems that have been left off this report either because they do not affect outward facing services, or have not yet been investigated enough to elaborate on them in detail.

4 Recent Work & Changes

4.1 New Systems

A new host, initially intended to be used for systems and network administration functions (monitoring, measurement, internal version control, configuration management, etc.) was added to the network in September. Many of these new tools had started to be deployed on a server temporarily borrowed from software development, and will be migrated to the new server in late October and November.

Although for now it is limited to administrative services, the function of this server will change over time as describe below in Future Work.

4.2 CVS to Git Migration

The CVS repository used for systems configurations, internal-only software development, and some document tracking has been migrated to a new GitLab instance. Barring any change in plans, access to this GitLab instance will remain staff-only. Public software development continues to be done via GitHub.

4.3 TLS Certificate Update

All of the dns-oarc.net web sites and OARC's jabber server were updated to use OARC's wildcard TLS certificate, replacing several self-signed certificates and site-specific certificates. The updates to OARC web sites were completed prior to the July 20th "first canary" release of Chrome, which would have invalidated OARC's older certificates. The Jabber server update was completed in September.

4.4 New Storage

A new drive shelf (JBOD) was added to Fs2 containing 24 12TB disks, the largest drives currently in service within OARC. The filesystem is configured using 20 of the drives in a ZFS RaidZ2 configuration with four spares to produce a 125TB filesystem.

This work was done in place of the planned upgrade of Fs2's existing ZFS filesystem (composed of 20 4TB disks) because there was insufficient space available on other file servers to move Fs2's data for the filesystem rebuild, and insufficient time available for a drive-by-drive resilvering of the entire filesystem.

4.5 Zone File Repository Rebuild

During the transition in the Systems Engineering role, it was necessary to reboot several systems in the process of recovering access, as they did not have the documented root password or remote console password set. `zfr.dns-oarc.net`, which hosts the Zone File Repository, was among these, however its network interfaces did not come back up after reboot. It was determined that the hardware had failed, and the service needed to be rebuilt elsewhere.

Unfortunately, the documentation for this service was incomplete, and it took many days of reading code and examining the old system (via java console) to discover all of the components that would need to be rebuilt. Along the way it was discovered that key data required by the service (the database containing the FTP access credentials and download attempt history, as well as the Subversion repository used for VCS access to the root zone archive) were not properly backed up, and these needed to be rebuilt or recovered via remote hands.

This process of recovery also revealed that several zones had been failing download for anywhere from weeks to years, and that other zones had been failing the stage that converted them to a canonical form, due to memory exhaustion.

The ZFR service has since been rebuilt on a virtual machine with double the RAM of the original physical host (formerly 4G, now 8G), and a new capability to obtain and archive zones via zone transfer has been added. Some of the larger holes in the archive have been plugged thanks to the contributions of James Mitchell at PCH and Roger Murray at IIS. At this time we are still missing the AERO and ASIA zones as of May 2016, the NET zone occasionally fails conversion to canonical form due to memory exhaustion, and the COM zone consistently fails this step. There are several other zones which have historically failed the conversion to canonical form, and a project to retroactively convert these zones will be taken up at a later date, likely after we have fixed the memory issues for COM and NET.

4.6 DNSSEC Validation Changes

In September, the DNS Privacy resolver testbed was updated to enable DNSSEC validation. OARC's other DNS recursive resolvers—the Open DNSSEC Validating Resolver (ODVR) service and OARC's internal recursive resolvers—were updated to enable RFC5011 automated updates in advance of the root key roll this week.

5 Future Work

5.1 The Do-Over

In order to address OARC's ongoing issues with legacy application architectures, poorly operating manually-compiled applications, and poor operating system distribution choice, we will be embarking on projects to containerize as many applications as possible, reimplement internal applications where necessary, and reinstall the operating systems on all OARC servers. Although these are actually three separate projects, they are inextricably linked and will need to be pursued concurrently.

Because OARC already has a fairly high operational load, the projects described below are expected to take a year or more to complete.

5.1.1 Reimplementation

There are several applications which are using years-old configurations which have been copied, unmodified, as they were originally implemented, from system to system over the years. These are in need of an update to more modern software, and configuration techniques. In many cases the current configurations are broken enough that it makes more sense to start over with a clean slate than to try to understand or update the original configuration.

Some examples:

- the current authoritative DNS infrastructure still implements DNSSEC signing using shell scripts written before the major DNS implementations did their own signature or key management
- the mail server is using an ancient sendmail configuration, is failing to filter most spam, and has several milters that fail to run
- dozens of cron jobs send mail even on success, making the target mailbox useless for finding actual problems (over 900 messages per week)
- the internal wiki configuration and library dependency tree was never completed or tested, resulting in several useful features (e.g. image upload and display) failing to operate
- hundreds of processes around the network cannot be operated using the standard INIT service tools

Reimplementation priority will be determined by balancing several factors including: squeakiest wheel; easily reimplemented applications; and, last remaining legacy apps on a host.

5.1.2 Containerization

Currently, OARC's systems are roughly divided into single service hosts (e.g. the capture or analysis servers) and multi-service hosts (mail, jabber, web site, portal, etc.). The multi-service hosts are in turn divided imperfectly into internal and external service delivery. All applications are installed and running on the OS on bare metal.

Most of the services on the multi-service hosts are prime candidates to be moved into Linux containers. This would improve security by isolating processes, simplify the upgrade process for most applications, and eliminate the occasional dependency conflict that we encounter.

Taking this one more step, and converting all multi-service hosts into a single containerized cluster (e.g. Kubernetes) would also allow us to make most of these services recover automatically from software and hardware failures, and even auto-scale capacity to a small degree.

Due to budget constraints we cannot simply build a cluster and migrate all applications to it. Instead, we'll be taking the new administrative host and making it the first node in a future Services Cluster. Applications will be containerized as they are reimplemented, and moved to this host. As hosts are vacated of all legacy applications, they will be added to the cluster which will then be allowed to rebalance load including the additional server.

If any applications are encountered which somehow defy containerization they will be moved into virtual machines.

5.1.3 Reinstallation

The experiment of converting the infrastructure to using Devuan in place of Debian, mentioned briefly in the last Systems Engineering update, has not been a success. Unfortunately, there does not appear to be a tested process for backing out of the conversion from Debian to Devuan. Rather than be the guinea pig testing back-out procedures, we will be reinstalling all systems currently running Devuan 1.x (based on Debian 8.x) with Debian 9.

Single-service hosts will be scheduled for reinstallation as the operational calendar allows. Each multi-service host will be reinstalled as the last of its services are migrated off the host onto the container cluster.

5.2 Configuration Management

Beginning with simple services, such as user management and logging, and later proceeding to more complex applications, OARC is going to begin the implementation of SaltStack to maintain systems configuration.

5.3 Log Consolidation

An ELK stack (Elastic, Logstash, Kibana) will be deployed to consolidate logging output from all applications network-wide into a single, searchable, graphable datastore.

5.4 Speculative Projects

The projects described below are simply ideas under consideration. They are not completely thought out, do not have budgets, or schedules, and in some cases are not even complete ideas.

5.4.1 File Server Clustering

The current architecture for OARC's large data storage is a group of individual file servers, each with one or more large redundant filesystems, each exporting those filesystems via NFS to a set of analysis servers. In addition to the data redundancy provided by the filesystems, each data set is stored on more than one file server in order to add cross-chassis redundancy of data, and in one case (Fs1 and Fs6) offsite redundancy.

The usual process for hardware refresh on these file servers that once a year, the filesystem made up of the oldest set of disks is chosen and:

- its data is moved to a different file server
- the drives are replaced

- the filesystem is rebuilt on the new drives
- the data is moved back

This arrangement has several scalability issues.

The rate of growth of OARC's data does not appear to be in line with whatever passes for Moore's Law in commodity disk drives. This forces us to periodically add a new file server or drive shelf. Because our budget doesn't allow us to upgrade the drives on multiple file servers per year (drives for the FS2 upgrade this year cost over US$10,000) this means that the number of years that any individual drive is expected to remain in service is increasing at a steady, if slow, rate. Maintaining the same maximum service lifetime as drive shelves are added would mean that budget increases would periodically involve significant increases rather than a steady increase year over year.

The inter-chassis data duplication is necessary, but operationally expensive. Because the file servers are not upgraded in pairs, the available storage on any individual server is unlikely to match the available storage on another server. Therefore, it is not efficient use of storage to simply have pairs of servers with identical data. Instead, there's a complex task of copying and moving data to maintain duplication and best-fit on multiple servers. As individual servers fill up, data has to be moved off of them onto newer (emptier) filesystems in order to make room for new data sets to be given a second copy. Additionally, verification of whether copies remain identical is a slow, resource intensive process.

As the number of file servers increases, so does the number of NFS mounts on analysis servers. This, combined with the need to periodically rebalance data sets, means that the paths to data sets researchers are using are moving targets.

With such large file systems on each file server, the effects of any single point of failure (e.g. a dead CPU) are significant.

And finally, the architecture is wasteful of space in two separate ways:

1. The combination of intra-chassis and inter-chassis duplication is required by the architecture, but wasteful, as it results in more copies of any individual block of data than is strictly necessary.
2. Because the data sets are large, best-fit organization of data still results in large chunks of unused space on each filesystem.

One obvious fix for all of the above problems is to convert the individual filesystems on individual chassis into a single large clustered filesystem.

- Clustering would reduce the total space usage by making all duplication cross-chassis, eliminating one of the dimensions of data redundancy that is currently necessary.
- Clustering would also eliminate the operational load required to manage cross-chassis redundancy, and simplify researchers' work by giving a single directory tree for all datasets.
- Over time, through regular hardware refresh cycles, the dependence on having large volumes of data on individual servers could be reduced by replacing the few large file servers with many smaller ones, minimising the impact of individual hardware failures.
- Drives could be replaced in a fixed schedule, with the annual increase in budget growing slowly and steadily, rather than in a large step function which would currently be required by a fixed schedule.

There are budgetary concerns with converting to a cluster, however.

It would be necessary to build a viable cluster large enough to accomodate the largest filesystem in the current deployment before any data migration could begin, with the remaining chassis individually converted over to the cluster as they are emptied. At present, this would be a significant increase in our annual capital budget. In theory this would be a short-term increase, resulting in capital savings further down the road, but that is not a given; the cost of enough small chassis to duplicate the service of one large chassis is currently higher than the single chassis. Even assuming there would be a return on investment, it's not clear how far in the future that would be, and that calculation, including potential savings in operational expenses, will need to be made before embarking on a project like this.

5.4.2 ELK As a Query Analysis Engine

Elastic Search is a highly efficient full text search and analytics engine. It's capable of quickly indexing very large data volumes, and of retriving billions of documents in extremely short (sub-second) periods of time.

It seems reasonable to consider the possility that PCAP or C-DNS data could be converted to JSON, indexed by Elastic Search, and then be made available for direct query via Elastic's API or for visualization via Kibana. This could significantly reduce the effort required to answer many questions asked of OARC's datasets.

Unfortunately, keeping both original PCAPs and a text conversion of all of OARC's DITL and other PCAP datasets could as much as double the current storage requirements.

However, this still seems like an experiement worth trying, as it could be one option among many that might improve data access times for simple research queries.

6 Conclusion

Although OARC's systems and network currently contain many challenges, there is a huge opportunity to improve services, reduce operational overhead, and more efficiently utilize resources. The next year should bring interesting changes.

Author: Matthew Pounsett <matt@dns-oarc.net>

Date: 2018-10-04