DNS-OARC Systems Update

DNS-OARC Workshop

Amsterdam, NL

May 09, 2015

DNS-OARC Services

- To recap, there are a number of services DNS-OARC provides:
  - Access to data archives (mostly DITL captures)
  - Mailings lists, websites, accounts, database servers, ODVR, secure jabber, don’t-probe, kitchen sinks, etc.,
  - Data collections (DITL, DNS-OARC testing services logs)
  - Dedicated capture systems to receive and process imported data types (DSC, ODVR and DITL data)
  - Various collections of support systems, like DNS, out-of-band management and network switches
  - Monitoring systems such as TLDMon, ZFR and DSC
  - And in one corner, custodianship of AS112 website.

System Status

- All looking good as of today
- All old servers have been retired, migrations have been completed
  - No DNS-OARC server systems older than 2013 are in operation
  - This does not count the DNS lab systems, as venerable as they are, which are still running well
- DITL systems partially rely on some older DNS lab systems, but are really occasional-use so risk is minimal

Data Archives Status

- All servers using ZFS, except fs1 which uses XFS as filesystem
- fs2, 15TB of 17TB used
- fs3, 21TB of 22TB used
- fs4, 42TB of 45TB used
- fs5, 5TB of 42TB used
- Thump1, 4.5TB of 14TB used
- Thump2, 0TB of 14TB used
- Thump3, 7TB of 13TB used (mirror of fs3 log holdings)
- Thump4, 14TB of 15TB used
- Total: 108.5TB used, 140TB total possible
- Access to that data available via analysis servers!
- Fs1 holds a copy of *all* data above, 110.5TB of 129TB used
- To be pedantic, 219TB used, of 281TB total capacity

New Data Arrivals

- DITL 2015!!
- SIE data
  - Thanks to Farsight for their contribution of a large data server (fs5) with old SIE data, 50% recovered
- RSSAC-002 archives
  - A, J, K and L root servers supplying this data - now mirrored by DNS-OARC locally.
- Long-term AS112 queries, 2-week durations of collections
  - Including data for before, during and after the surprise RFC6304bis (RFC 7534) deployment

DITL 2015

- Thank you to participants for making DITL 2015 a success!
  - A number of AS112 participants this time around
  - Some in.addr-arpa and ip6 included
  - Makes this the most diverse DITL collection year ever
  - Special thanks to OttIX for providing early testing
- 6.4TB of raw data, converted to 3.0TB of clean data, 9.4TB total
  - Note: There will be a second run performed to verify all RAW/ data is ‘cleaned’, since quite a bit of the dataset is new to our tools
- Available now to all members in good-standing on all analysis servers
  - All analysis servers were upgraded to 10Gb/s in mid-April 2015

DITL 2015 Post-processing

- But this posed a slight problem:
  - No room to process the uploaded data on fs2 with 2 dying disks
- Ix2 was picked as the spool for post-processing since it had disk space, but needed 10Gb/s NIC
- Once installed, the final turn-up for 10Gb/s took place April 22, resulting in the following paths:
  - DITL-20150413/RAW/ on fs2
  - DITL-20150413/CLEAN/ on ix2
- Processing roughly took about 2 weeks, with fs1 doing the same processing and finishing in 3.5 days
- For more information:
  - https://www.dns-oarc.net/node/350

DITL 2015 Network Stats

- Typical sub-500pps traffic rose to average 14500pps
- Average bandwidth settled on 120Mb/s
- Connections peaked around 1000/s
- All this was detected through Netflow v9 traffic data export, both on IPv6 and IPv4 connections, using nfsen and nfdump
  - Netflow is very cool for other things…

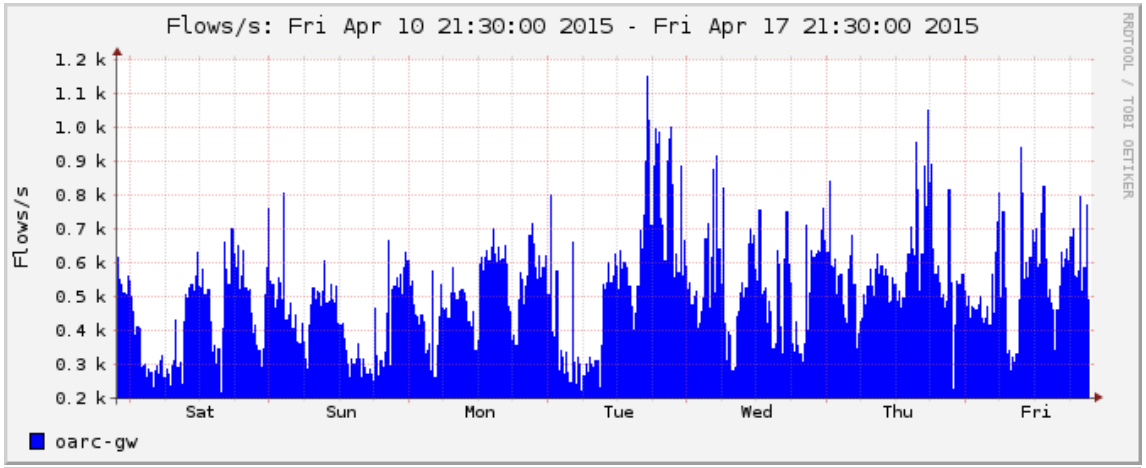
Flows/s: Fri Apr 10 21:30:00 2015 - Fri Apr 17 21:30:00 2015
Flows/s
1.2 k
1.1 k
1.0 k
0.9 k
0.8 k
0.7 k
0.6 k
0.5 k
0.4 k
0.3 k
0.2 k
Sat
Sun
Mon
Tue
Wed
Thu
Fri
oarc-gw
RRDTOOL / TOBI OETIKER

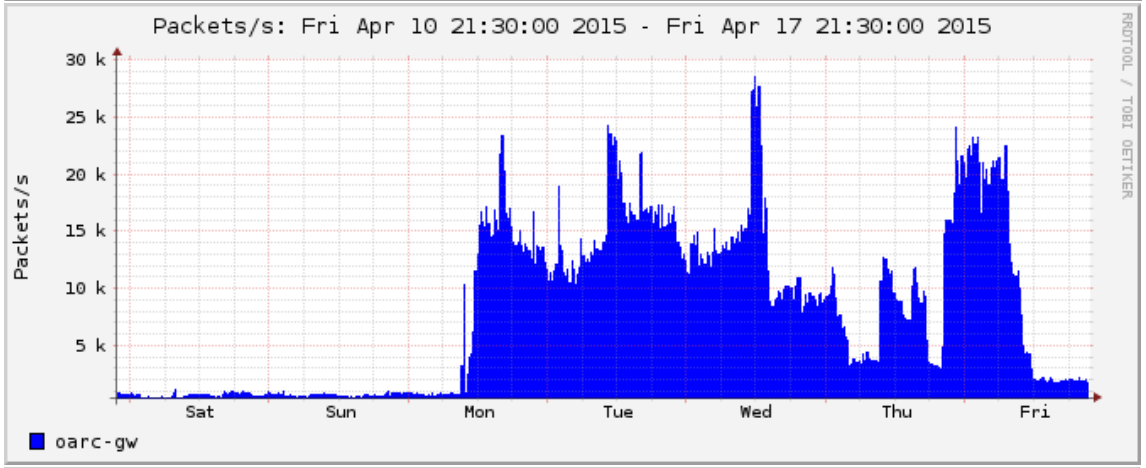
Packets/s: Fri Apr 10 21:30:00 2015 - Fri Apr 17 21:30:00 2015
Packets/s
30 k
25 k
20 k
15 k
10 k
5 k
Sat
Sun
Mon
Tue
Wed
Thu
Fri
oarc-gw
RRDTOOL / TOBI OETIKER

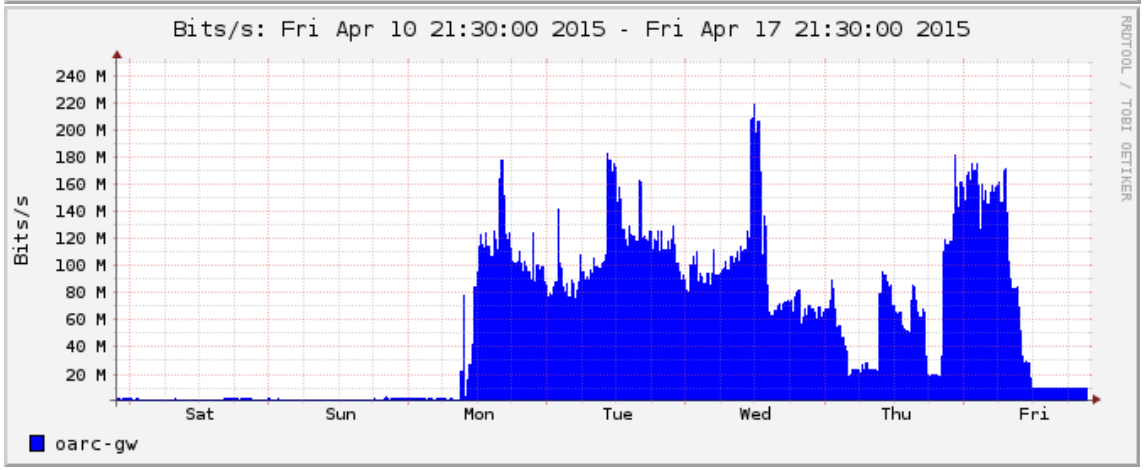
Bits/s: Fri Apr 10 21:30:00 2015 - Fri Apr 17 21:30:00 2015
Bits/s
240 M
220 M
200 M
180 M
160 M
140 M
120 M
100 M
80 M
60 M
40 M
20 M
Sat
Sun
Mon
Tue
Wed
Thu
Fri
oarc-gw
RRDTOOL / TOBI OETIKER

DITL Timeline and Volumes

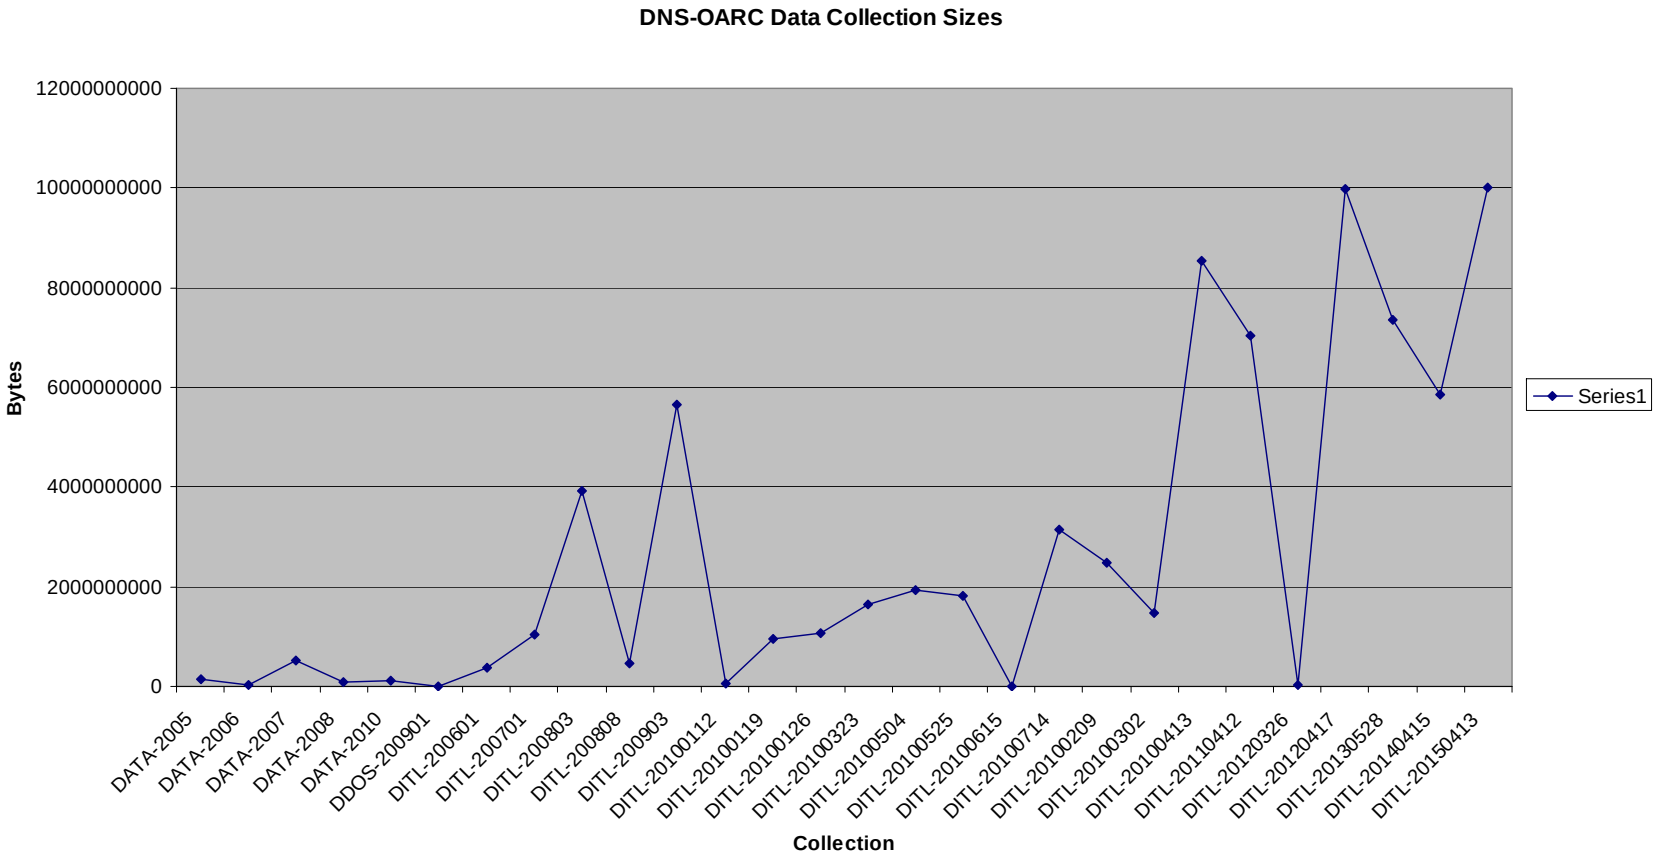

A Note to Researchers

- We welcome researchers, but:
  - Any output of your research work must be reviewed with DNS-OARC as per data access agreement
- This ensures that DNS-OARC plays its part in knowing who publishes papers giving it the opportunity to create pointers to that output as part of a library
- DNS-OARC can also demonstrate to the board that data collected is being used and has value
- It also reassures data providers that their data is in fact of use to the community, and not just being collected for collecting's sake.

DITL 2016

- Some proposed changes are coming to the way the next collection will be made
  - HPN-SSH, a high-performance version of SSH used by the NREN's & Science to speed up data transfers using SSH
    - http://www.psc.edu/index.php/hpn-ssh
    - Recommended reading: http://fasterdata.es.net/
    - (For the adventurous: Would tuning systems increase or decrease DNS performance?)
  - Relocating DITL uploads in another country, along with the whole DITL data catalogue
  - Having separate upload accounts for each query group (AS112, in.addr-arpa, etc)
  - Switching to a more efficient compression tool, like xz which has consequences for the entire DITL analysis tool chain and researchers
- Watch for announcement of next DITL test and DITL dates after the New Year.

Backup Strategy

- Backups are routinely performed on all systems using BackupPC
  - Except DITL and DSC data which are duplicated to fs1, in turn to be duplicated offsite
- Retention between 30-45 days, including full and incremental versions
  - DITL and DSC, forever
- To date there are no holes with regards to data backup.

Data Access

- Data access is made through 4 dedicated systems: an1, an3 (both FreeBSD) and an2, an4 (both Linux) via SSH
- Members and research users are encouraged to make use of limited resources in a responsible manner
  - The alternative is to enforce a quota system and forcibly limit access to CPU and memory
- All these systems, including the file servers now have jumbo frames enabled (9000 byte MTU) for better NFS throughput
- Various systems now have 10Gb/s interconnects between each other to help facilitate speedier data access, again 9000 byte MTU
  - Fs2, fs5, ix2, an1, an2, an3 and an4 (stay tuned for more!)
  - Note that the bottlenecks will likely be system I/O

DNS-OARC Portal Status

- Improvements mainly to help back office operations more efficient
  - More org info, such as postal/mailing address
  - Signatory, admin and billing contacts (new)
  - Shipping/mail contacts (re-implemented)
  - Database tools (member/user flat-file exporting - new)
  - Data Providers info/updates (new)
- “My DNS-OARC” web page no longer redirects after login to the defunct incidents pages
- Member requests for improvements welcome.

Portal Reminders

- Members are automatically subscribed to the closed members mailing list
  - Please note: You must use the same email address to post to the list as entered in the Portal or your message will silently bounce and be discarded
- For those who ask for a secure Jabber account (@dns-oarc.net), they are automatically created nightly on the Jabber server
  - This is a secure jabber server cluster, channel-wise, but does permit loose associations with other Jabber servers on the Internet
- Need a password reset? Where's your PGP key??

Other Interesting Bits

- DANE records have been deployed for certain websites and oddballs using SSL certificates
- OpenDKIM and OpenDMARC has been enabled, even for mailing lists, but in a relaxed configuration to simply see how all that works
- A secondary DNS & MX for DNS-OARC exists in Sweden in case of access outages at DNS-OARC HQ.
- Mailbots sending occasional reminders to DNS operators about DNS-OARC services will become the norm to help keep the community informed.

Top DNS-OARC Services

- If looking strictly at web visits, what tools are hot over a typical 8-week period?
  1) DNS Entropy Test
  2) Reply Size Test
  3) Port Test
  4) Dnscap
  5) TLDMon
- The rest of the visits are spread over informative topics, such as minimizing DDOS, ODVR, DITL, etc.

Future

- Fs6, a clone of fs1, not located on the west coast of the US.
  - Any suggestions for possible co-lo locations welcome and prospective hosts please do drop us a line at admin@dns-oarc.net
- Spreading the services out to other places
  - Hosting requirements document had been developed and will be used to inform relationships and expectations
  - This includes a possible TLDmon deployment to a number of regions, using 1U servers according to standard specs
- Encourage greater visibility within the Internet2 (research) community for more mutual research opportunities (increase the ‘R’ in DNS-OARC)
- Relocating upload and analysis elsewhere
- Upgrade of jabber server software to support IPv6 connections for those who crave such things
- With new disk space, we may re-process the whole archive and update the DITL summary pages accordingly.
  - Release of DITL 2005-2015: The Greatest Pcaps

Supplemental Slides: Monitoring Status

TLDmon, ODVR, ZFR, DSC and Friends

TLDMon

- A set of Nagios plugins monitoring TLD DNS server states
- 573 TLDs as of May 1, 2014
- 723 as of Sept 23, 2014
- 790 as of Jan 8, 2015
- 931 as of May 8, 2015
- As the root zone grows, it will be interesting to see how it and this tool scale.
- Most email notifications related to monitoring of TLDs go to ICANN
  - If you (members that is) would like to receive notices directly for your TLD, please contact admin@dns-oarc.net
- https//tldmon.dns-oarc.net/nagios/

TLDMon Expansion

- Currently prototyping another instance of TLDmon to gain a secondary view of the global DNS from another perspective, and gather yet more data
  - We are not soliciting the hosting of other TLDmon servers quite yet as the value in doing so needs to be better understood.
  - TLDmon's 'measurements' are only as effective as the closest anycast server, how does this affect them and can that anycast server always be uniquely identified?
  - TLDmon EDNS testing (but is not an in-depth EDNS0 compliance tester)
- Using tldmon.cc.dns-oarc.net as a standard naming convention (where cc is the country code where this instance is located)
- Test version is tldmon.ca.dns-oarc.net

ODVR

- Open DNSSEC Validating Resolver
- ODVR gets quite a lot of use
  - It is highly probable that the *use* is more like *abuse* rather than testing DNSSEC
- Access to service may be restricted to those who require it, instead of being as open at it is now, however:
  - There is a wealth of data there to be mined and it shouldn't be let go to waste
- Some thought is being given to expand the service to test other DNS servers, since this open resolver is quite popular
- Nonetheless, blatant abusers have been blocked if the system or service is affected.
- If the original purpose of ODVR has been fulfilled, it may then:
  - Be shutdown and the work declared concluded
  - Re-purposed into testing a suite of different resolver types
  - Mmmmmm….data……
- https://www.dns-oarc.net/oarc/services/odvr

ODVR Stats

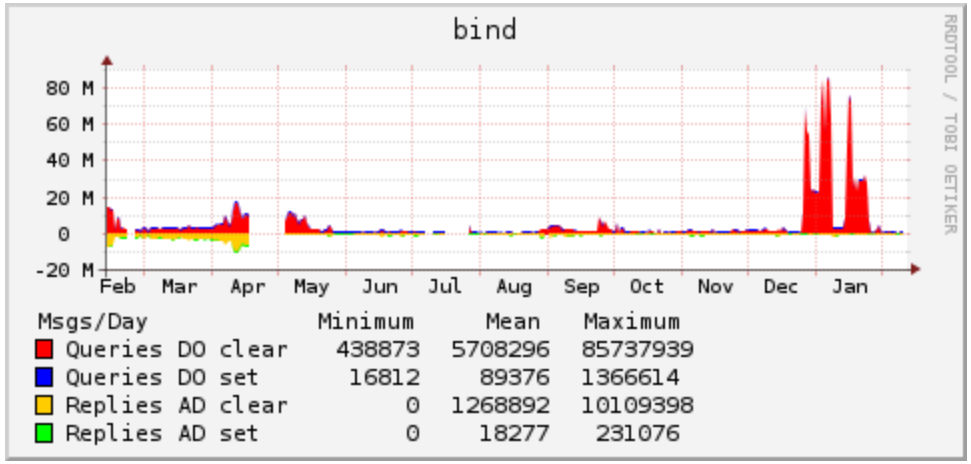
bind
80 M
60 M
40 M
20 M
0
-20 M
Feb Mar Apr May Jun Jul Aug Sep Oct Nov Dec Jan
RRDTOOL / TOBI OETIKER
Msgs/Day Minimum Mean Maximum
Queries DO clear 438873 5708296 85737939
Queries DO set 16812 89376 1366614
Replies AD clear 0 1268892 10109398
Replies AD set 0 18277 231076

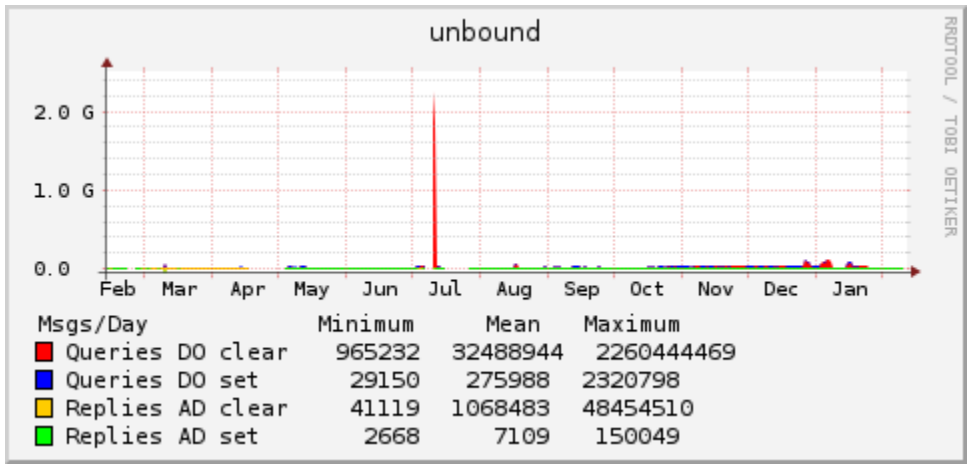
unbound
2.0 G
1.0 G
0.0
Feb Mar Apr May Jun Jul Aug Sep Oct Nov Dec Jan
RRDTOOL / TOBI OETIKER
Msgs/Day Minimum Mean Maximum
Queries DO clear 965232 32488944 2260444469
Queries DO set 29150 275988 2320798
Replies AD clear 41119 1068483 48454510
Replies AD set 2668 7109 150049

DSC as Software Package

- DNS Stats Collector
- Some renewed interest recently for DSC
- However, we're running into the yummy/yasty/rpm-y crowd
  - "Do you have an rpm for that?"
  - There already is an independent Debian package that could be used in a pinch
- May consider doing binary distributions but this may be more work than is worth.
  - Perhaps someone in the community would care to contribute?
- The subversion repository is at DNS-OARC

DSC as an DNS-OARC Service

- As an DNS-OARC service, continues to operate normally
- ISC is the latest returning contributor (AS112 node)
- Uploading and processing of XMLs has benefitted from performance gains using jumbo frames on the analysis network
  - But during the summer of 2014 we switched to a local machine to process them directly instead of over NFS
- As of Sept 10, 2014 a special binary now handles processing of the 5 minute interval XML files
  - 5 minutes of DSC XMLs now processed in 0.7s, was 2s for the same amount before, and before the jumbo frames upgrade, 3.6s
- The unintended consequence is email notifications coming from the system stating that it doesn't have anything to process…
- The DSC directory tree has also been cleaned-up and duplicate trees removed, with all data now in one place per data provider
- View more via the Portal
- https://www.dns-oarc/net/tools/dsc

DSC Member Contributions

- We encourage members to consider a restart of DSC uploads, although in order of preference, via rsync of post-processed XMLs followed by the usual upload method and having us do the processing
- The development of newer variants of DSC, such as DSCng and Hedgehog has however created a barrier in sharing DSC data with DNS-OARC (and by extension, others)
- A potential DSC project would be to create ways that data stored in DSC descendants can be exported to facilitate information exchange using a standard format
  - Most likely just post-processed DSC XMLs

dnscap

- A tcpdump-like tool for DNS packet captures
- Usually used for DITL captures as part of the larger ditl-*.tar.gz capture suite
  - Good for collecting packets on one machine with many physical interfaces
- Periodic updates
- Newest versions and contributions available via gitub
- https://dnscap.dns-oarc.net/

ZFR – Zone File Repository

- ZFR records changes in a select bunch of zones monitored
  - Recorded in subversion repository
- Most important of these is the root zone portion.
- Has served the community well
- Can be expanded to track more TLDs and has the capacity to do so
  - We would welcome more additions to the archive as this resource is severely under-nourished and utilized with a very large potential to be useful
- Data accessible via dedicated web portal or via the analysis servers
  - Should this be for members only?
- Suggestions of which zones to add are welcome, contact admin@dns-oarc.net
- Same for those TLDs who wish to use us to monitor on their behalf to act as a double blind.
- https://www.dns-oarc.net/oarc/data/zfr

How Big is the Root Zone?

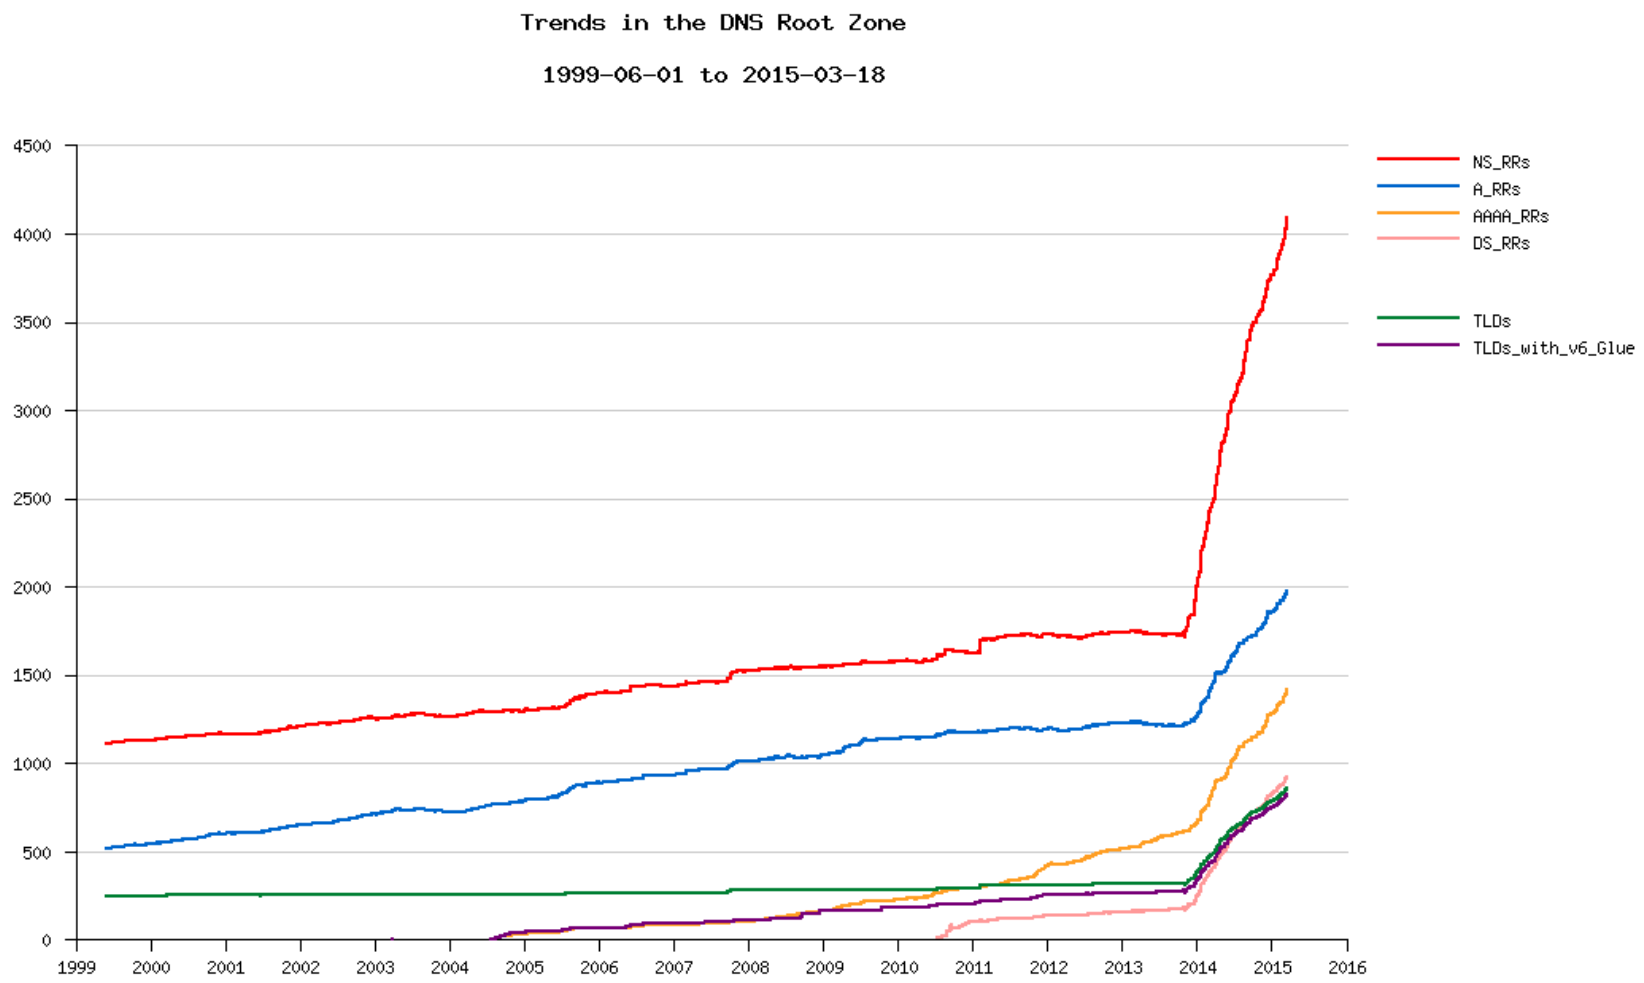

Don’t Probe Database

- DNS-OARC hosts a no-scan/blacklist of DNS servers and networks which
  - DNS researchers are encouraged to obtain a copy of; and,
  - DNS operators are encouraged to request having their DNS servers listed to avoid scans from researchers.
- https://dontprobe.dns-oarc.net/

Reply Size Test

- A very popular service to test the DNS reply size of DNS servers via dig
- Dual-stack, does respond to both IPv6 and IPv4 addresses
- We do maintain logs of the use of the tool, which may be of interest to researchers
- Thanks to Mark Andrews for pointing out a shortcoming in the detection and test code and Duane Wessels for the fix, now EDNS compliant
- https://www.dns-oarc.net/oarc/services/replysizetest

__END__