DNS-OARC Systems Update

DNS-OARC Workshop

San Juan

March 8, 2018

System Status

- Services, systems and data archives all operating normally
- Fremont, California, USA has:
  - All services and main systems, including data
  - Three networks: Public with 1Gb/s upstream for transit and peering, private and private high-capacity 10Gb/s
- ‘DR’ site in Ottawa, Canada has:
  - an3 and fs6 (data backup) in case main systems are unavailable
  - Sb (system backup) and ns2 (secondary DNS)
- Site in Sweden has:
  - Ns3 (tldmon-se, backup MX and secondary DNS)
  - Dev, development platform
- DNSLab hardware is effectively shutdown, no use in 3 years

Portal Reminders

- Members are automatically subscribed or unsubscribed to the closed members mailing list through the Portal *only*
  - 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
  - Subscriptions are exclusively handled through the Portal, not the mailing list interface
- For those who setup a secure Jabber account (@dns-oarc.net) in the Portal, they are automatically created nightly on the Jabber server
  - This is a secure, private jabber server and does *not* permit associations with other Jabber servers on the Internet

File Servers & Datasets

- Current total dataset size sitting at 231TB used of 247TB total, not including 2 duplicates of 125TB each, therefore 0.5PB total capacity:
  - Fs1: 125TB used, 129TB capacity
  - Fs2: 33TB used, 45TB capacity (upgrading to 90TB)
  - Fs3: 34TB used, 45TB capacity
  - Fs4: 55TB used, 90TB capacity
  - Fs5: 112TB used, 134TB capacity
  - Fs6: 125TB used, 129TB capacity (fs1 mirror)
- FreeBSD 11.1 using ZFS on fs2, fs3, fs4 and fs5 and Debian 8 using XFS on fs1 and fs6

Analysis Servers

- an1, an2, an3 and an4 are available.
  - Note: Expired members have their accounts deleted.
  - After 2 years unused accounts and data will be purged
- Not much analysis activity these days.
- An4 and ix2 were swapped to accommodate an urgent demand from a researcher for a limited time who required more locally attached disk capacity.
  - An4 with 64GB of RAM, 16TB of disk, E5-2620 2GHz
- All servers have been upgraded to Debian 8 for more current library support, with an eye to Devuan transition later
- **Note well**: No data, even derived data, may leave DNS-OARC systems without express written authorisation. Contact admin@dns-oarc.net first, *always*

Analysis Server Considerations

- Data ex-filtration prevention with respect to EU GDPR
- Looking ahead, a proposal to clone an4's hardware and offer exclusive analysis services
  - Dedicated to one research user, not shared
  - Limited duration
  - Wiped and rebuilt for next researcher
  - Stay tuned for a future Executive Director's member report for more information as this evolves
- Also considering attaching additional disk space to individual analysis servers over SAS external connection
- Shared scratch
  - Current analysis server fleet has small local disk footprint which has served users well
  - In HPC, common to see ephemeral, writable, large scratch shared between users
  - More cost effective than replacing existing CPU and memory rich systems
  - Make available over NFS via 10Gb/s analysis network

New Data Arrivals

- Long-term data acquisition:
  - AS112 queries from one “global” node continue
  - Annual Case Connection Zone DNS Transactions
- In past 12 months, major collections:
  - DATA 2017
  - DITL 2017
  - KSK 2017

DITL 2018 News

- Coming in April
- Test window March 14, 2018
- Last collection before impending EU GDPR
  - Various anonymisation tools suggested in the past for IP addresses, look for those if needed
- Capture servers have been replaced with more modern sleek versions, from a gang of 4 power hungry boxes to just 2 but same SSH host key and same IP addresses reused
  - Capture-c and capture-d are not available, check your hard-coded configs to ensure you use the round-robin hostname!

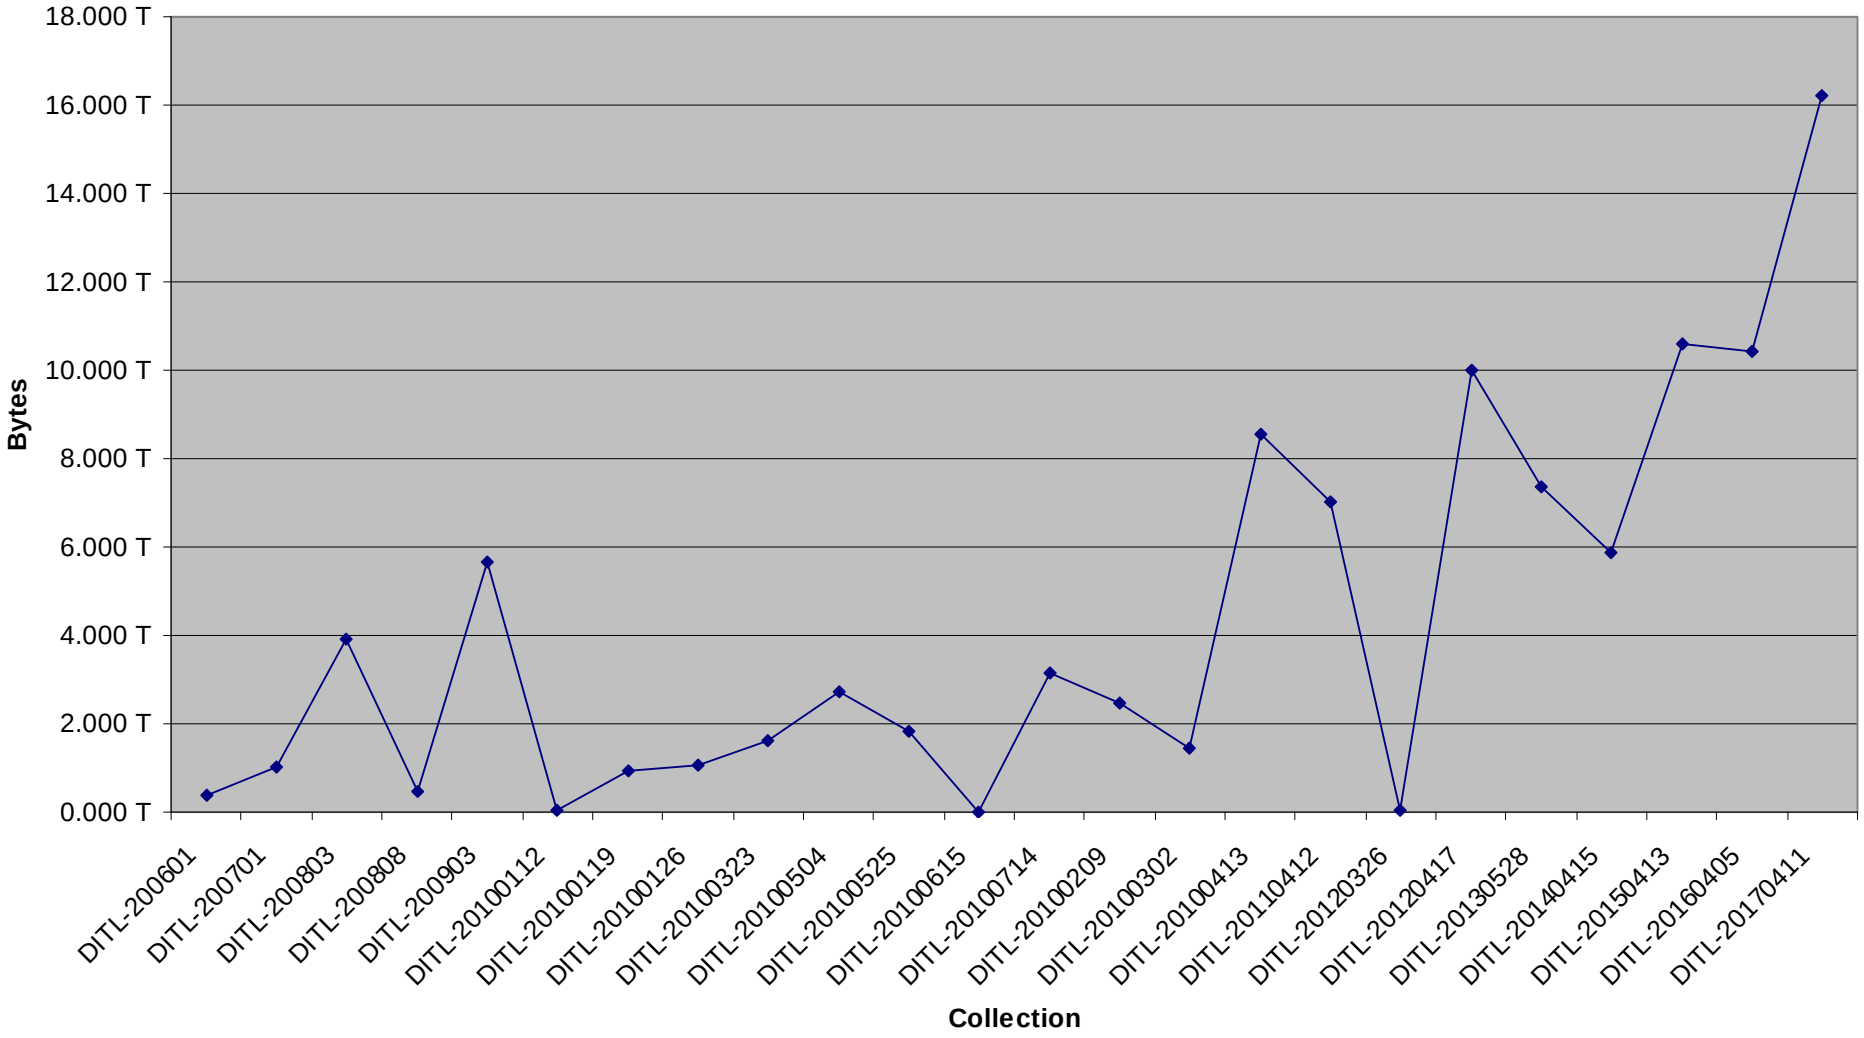
DNS-OARC DITL Data Collection Sizes
18.000 T
16.000 T
14.000 T
12.000 T
10.000 T
8.000 T
6.000 T
4.000 T
2.000 T
0.000 T
Bytes
DITL-200601
DITL-200701
DITL-200803
DITL-200808
DITL-200903
DITL-20100112
DITL-20100119
DITL-20100126
DITL-20100323
DITL-20100504
DITL-20100525
DITL-20100615
DITL-20100714
DITL-20100209
DITL-20100302
DITL-20100413
DITL-20110412
DITL-20120326
DITL-20120417
DITL-20130528
DITL-20140415
DITL-20150413
DITL-20160405
DITL-20170411
Collection

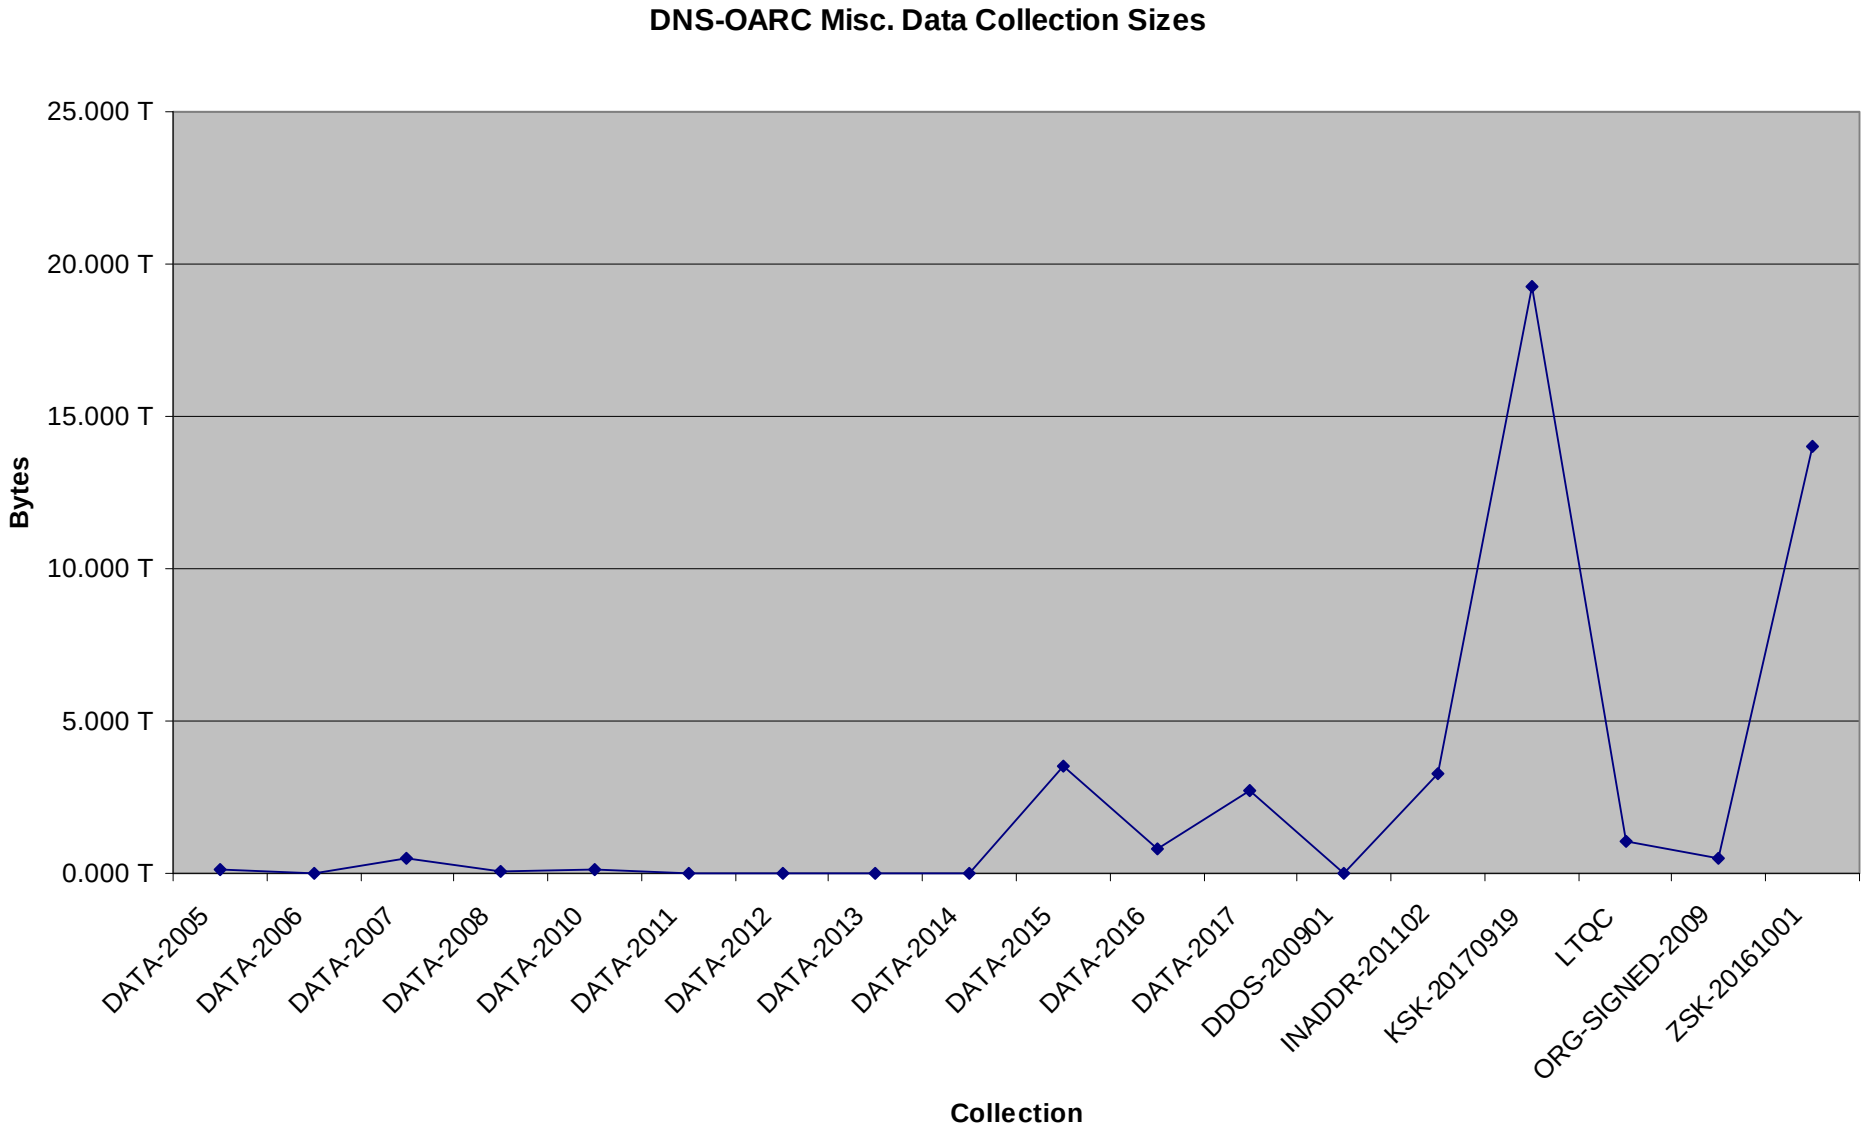
DNS-OARC Misc. Data Collection Sizes
25.000 T
20.000 T
15.000 T
10.000 T
5.000 T
0.000 T
Bytes
DATA-2005
DATA-2006
DATA-2007
DATA-2008
DATA-2010
DATA-2011
DATA-2012
DATA-2013
DATA-2014
DATA-2015
DATA-2016
DATA-2017
DDOS-200901
INADDR-201102
KSK-20170919
LTQC
ORG-SIGNED-2009
ZSK-20161001
Collection

DITL Data Formats

- PCAP is currently the *lingua franca* which researchers understand and expect to deal with in their current workflows
- There is a growing fragmentation of DNS data formats being suggested and presented to DNS-OARC in response to the perceived waste PCAP files contain
- Before proceeding with further fragmentation the community needs to consider a standard data format for representing captured DNS packets that will withstand the test of time, be accessible and best serve the needs of DNS dataset research.
- Speaking of competing formats, how about what kind of anonymisation of data is used and whether it matters?

KSK 2017 Collection

- First collection uploads started on Sept. 18, 2017 (most started as scheduled on the 19th), 3-4 days collected
  - 11.2TB raw, 8.0TB clean, 19.2TB total
- I-root holds the record for longest DITL-like upload (almost 5 weeks!)
- Initial processing completed Nov 14, 2017.
  - Further processing completed January 26, 2018 due to additional upload by F.root

KSK 2018

- Second KSK collection postponed due to rollover delay
- May happen during the proposed October 11, 2018 date
- DNS-OARC will be ready to receive data if called upon.
- Just use the same tools as in the usual DITL collection event

Other News for Data Archive

- Fs1 and fs6 backups are full
  - Data backups are local again between fs2/3/4/5 which will fill them faster
- May consider a complimentary backup scheme where each backup file server has a portion of datasets, bringing the copies from 2 to 1.
  - Older datasets kept remotely, newer copies kept locally
- Alternatively, fs5b could be upgraded to use 8TB disks, making it a 180TB archive machine
- Other possibilities include using the on-board SAS expanders to add more chassis
- There is a growing collection of idle, high-capacity disks (4TB) which would be useful in extending archiving and swing space options, as well as for scratch space
- Or keep only last 3 years of data and delete the rest.

TLDMon

- http://tldmon.dns-oarc.net/
- TLDMon sees 1543 TLDs as of Feb. 23, 2018
  - Down one from Sept 2017
- Experimental TLDMon-ca node in Canada
  - Sees 1790 zones combined (.arpa, etc.)
  - Nagios is doing well to handle that load so proves to be scalable for more TLD growth on main service
- TLDmon Nagios core upgraded to 4.x and works unmodified

ZFR

- ZFR still going strong
  - Some TLDs archived as well as the Root Zone Archive
  - Suggestions for other notable ones to track welcome
  - Root zone examples from each of 1993 and 1994 years now included
- A couple of zones have been migrated to ICANN's CZDS and away from traditional FTP downloads
- Integration with their API and ZFR is done
- /mnt/oarc-pool3/ZFR/ to find them

Root Zone Trend

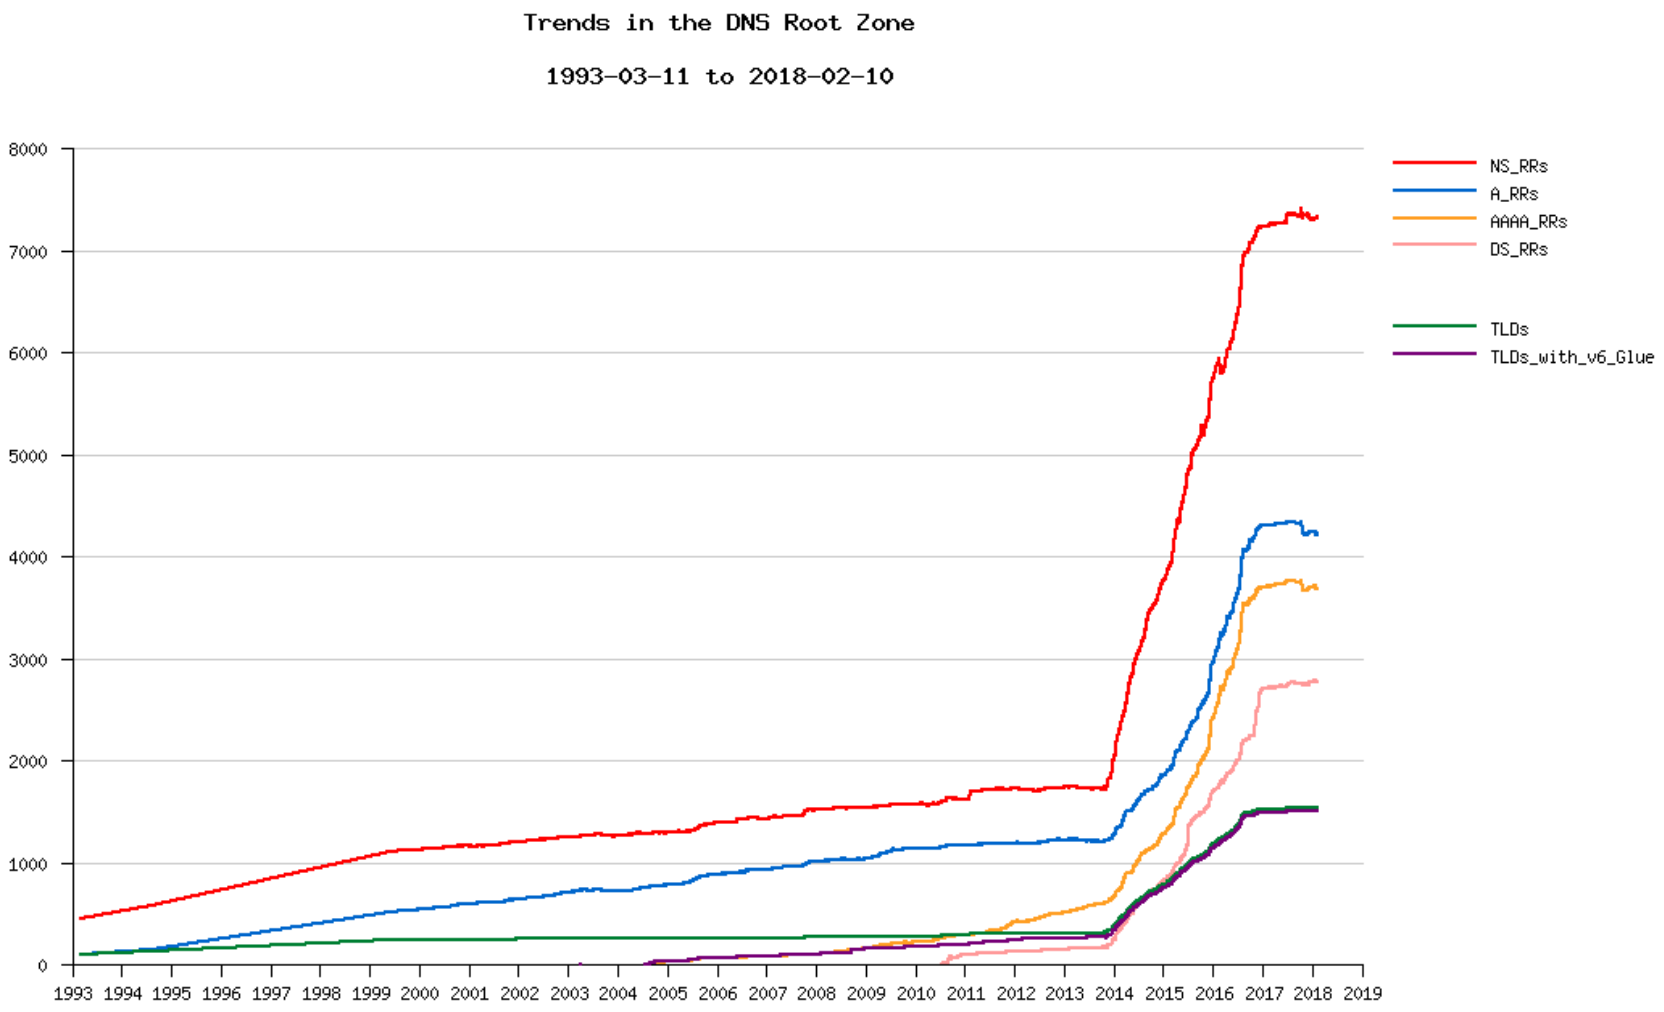

DS RRs to Date

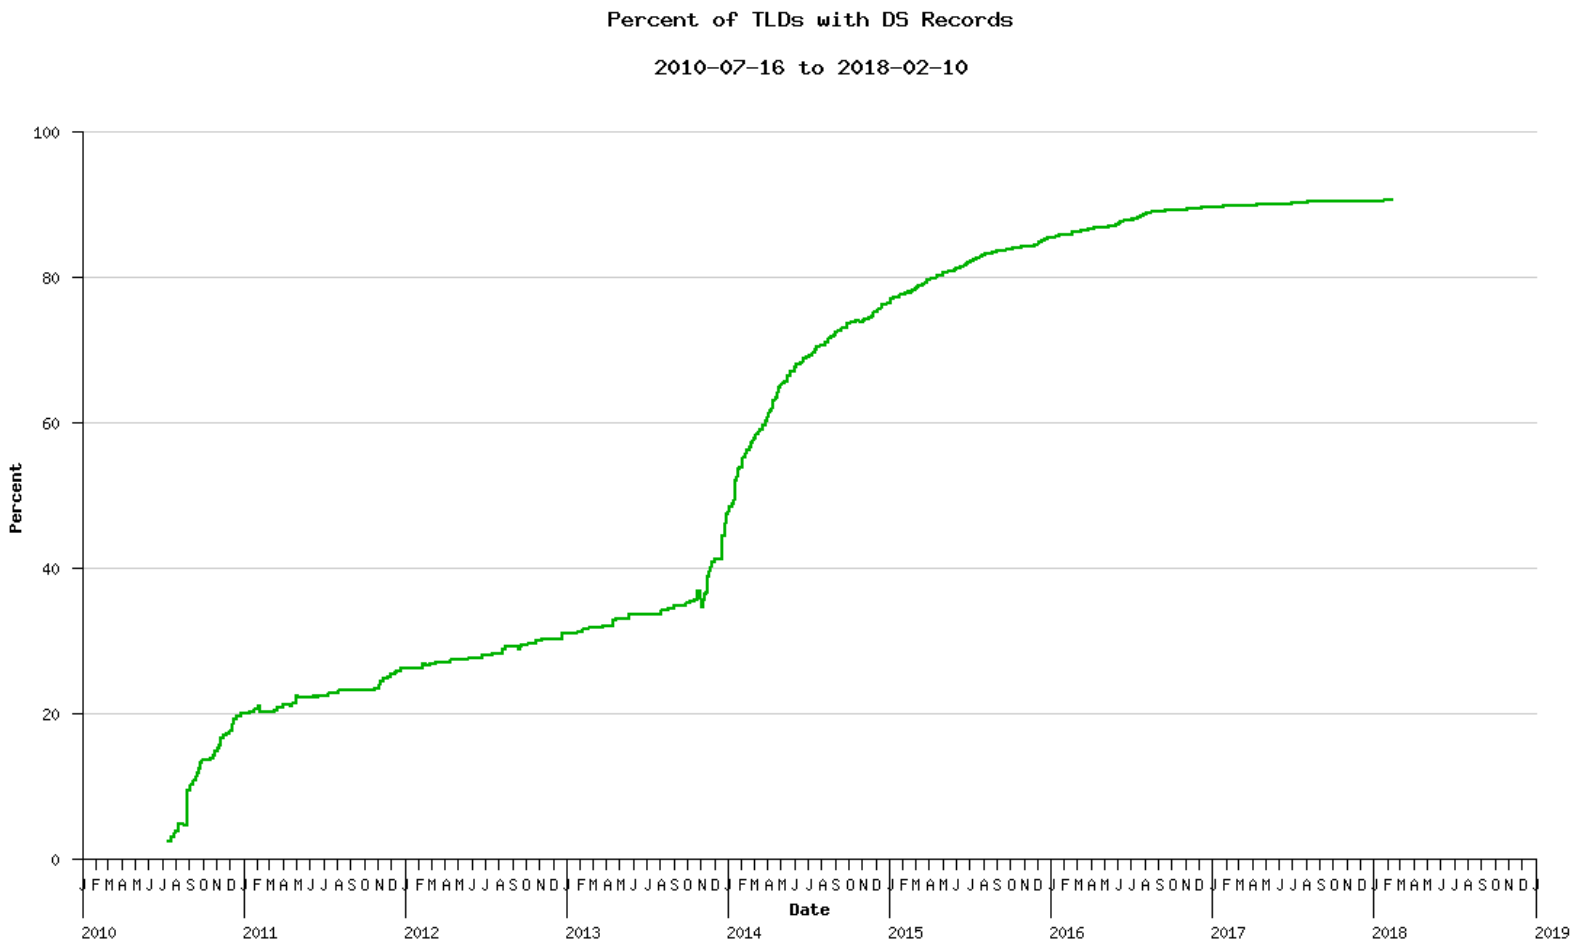

ODVR & TLS-DNS

- ODVR maintains a passing acquaintance with its DNSSEC test bed origins
  - It's really just an open resolver, with all that it entails
  - Updated trust anchors for KSK installed for both BIND and Unbound versions of ODVR
  - This may be a good source of info on the less than 2% who actually use this as a validating DNSSEC resolver
  - Would be nice to simply REFUSE queries not using DNSSEC
- Operates as a back-end to the TLS-DNS privacy service
  - Some abuse of this service has already occurred
- Logs and pcap traces preserve the data
  - 1.2TB of data so far. What is inside it?
- /mnt/oarc-pool3/odvr/ houses them

ODVR

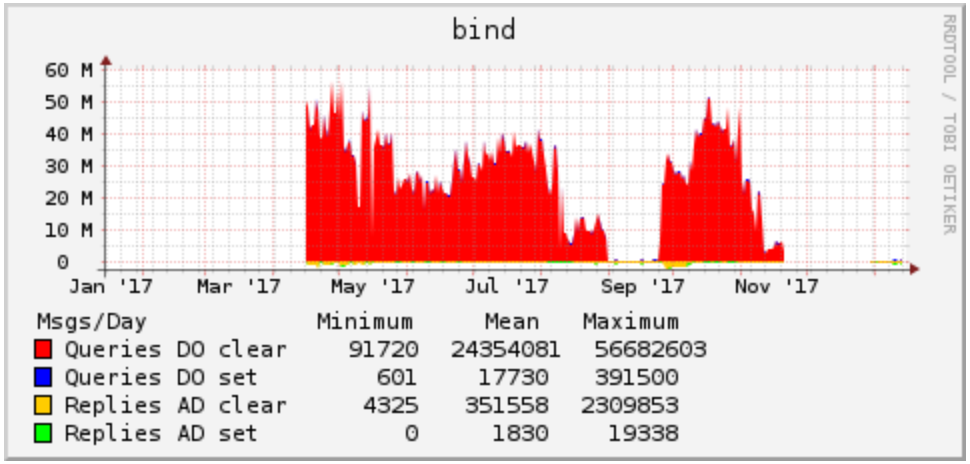
bind
60 M
50 M
40 M
30 M
20 M
10 M
0
Jan '17
Mar '17
May '17
Jul '17
Sep '17
Nov '17
RRDTOOL / TOBI OETIKER
Msgs/Day Minimum Mean Maximum
Queries DO clear 91720 24354081 56682603
Queries DO set 601 17730 391500
Replies AD clear 4325 351558 2309853
Replies AD set 0 1830 19338

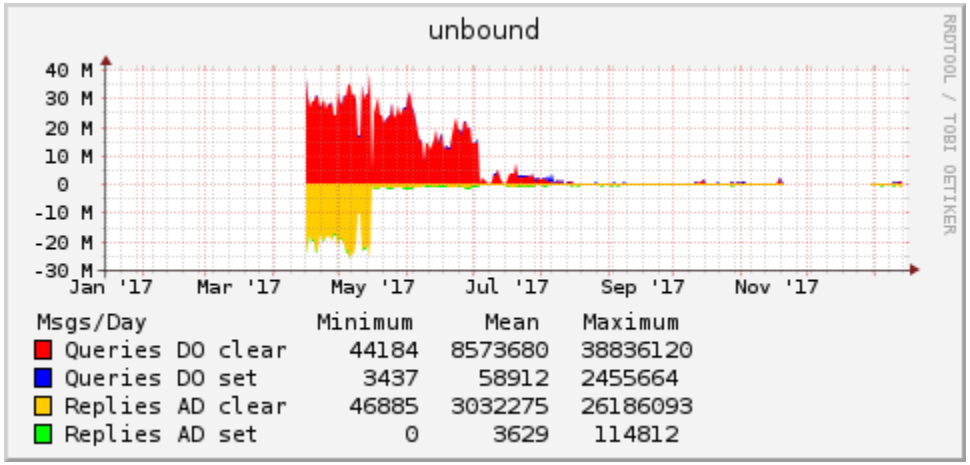
unbound
40 M
30 M
20 M
10 M
0
-10 M
-20 M
-30 M
Jan '17
Mar '17
May '17
Jul '17
Sep '17
Nov '17
RRDTOOL / TOBI OETIKER
Msgs/Day Minimum Mean Maximum
Queries DO clear 44184 8573680 38836120
Queries DO set 3437 58912 2455664
Replies AD clear 46885 3032275 26186093
Replies AD set 0 3629 114812

Future

- More capacity growth using 8TB disks or larger in fs1 and fs6 fileservers, or increase the number of fileservers to use dormant spares, perhaps as early as 2018
- Potentially offer to install and run AS112 nodes in certain locations.
- Indico 2.0.1 major upgrade for mid-2018
- Further consideration towards an analysis cluster
  - Hadoop and/or OpenMPI
- Data catalogue descriptions (ie., what is it?) for each collection
- Portal refresh – some call it plastic surgery, others just “restoration.”
- Dismantling the DNS Lab
- More data handling considerations
  - Does someone have the right to have their DNS queries forgotten?
  - Anonymity of the data: What type used and whether different types actually matter

Future II

- Consider studying various compression methods but there needs to be a balance between:
  - Transmission (minimize bandwidth use)
  - Fast decompress (saves time for researchers)
  - Minimize storage (save $$ on disk & power)
  - Reliability (very long-term archiving)
- Whatever future compression method is chosen, a long-term commitment must be made so it needs to be a wise decision.
- Some compression tools are available for parallel processing, such as those using OpenMPI

__END__