DNS-OARC Systems Update

DNS-OARC Workshop

Madrid

May 14, 2017

System Status

- Services, systems and data archives all operating normally in Fremont, California, USA
- Indico, works well but is becoming dated
- All analysis and fileservers now have uniform NFS mounts through the 10Gb/s network
  - This network is also used for inter-file server communications
- ‘DR’ site in Ottawa, Canada has:
  - an3 and fs6 (data backup) in case main systems are unavailable
  - Sb (system backups, truly offsite) and ns2 (secondary DNS, second-half of Jabber)
- Site in Sweden has:
  - Ns3 (tldmon-se, backup MX and secondary DNS)
  - Dev, development platform

File Servers

- Current total dataset size sitting at 167TB used of 247TB total, not including 2 duplicates of 129TB each, therefore 0.5PB total capacity:
  - Fs1: 120TB used, 129TB capacity
  - Fs2: 29TB used, 45TB capacity
  - Fs3: 34TB used, 45TB capacity
  - Fs4: 40TB used, 90TB capacity
  - Fs5: 64TB used, 67TB capacity
  - Fs6: 120TB used, 129TB capacity (fs1 mirror)
- Switchover to xz compression for the whole archive on hold due to long-term data reliability concerns with xz
- Some re-arranging of data on file servers yielded gains

Analysis Servers

- Both an1 and an3 have been converted into Linux systems from FreeBSD
  - Both systems are ready for research use
  - Ask admin@dns-oarc.net for any software required for system-wide use
- Aim is to achieve a uniform support footprint but more importantly something much more special
- Thus an1, an2, an3 and an4 are uniform in OS
- Stay tuned…

New Data Arrivals

- We are up-to-date in terms of mirroring all known publicly published RSSAC-002 metrics
  - All roots except G now collected
  - Issue surrounding firewalling of DNS-OARC address space by DISA is being investigated
- Long-term AS112 queries from one node continue, with contributions from ICANN and others in spurts

ZSK rollover in Oct. 2016

- Fs4 was upgraded from 45TB to 90TB using 8TB disks, in time for this collection, thank you to ICANN & Verisign
  - Processing and databases were also relocated here as a result
  - In time for…
- ….A query-only collection for Oct 1-4 2016 was completed
- This was for tracking the effects of the ZSK rollover in the root
- 7.4TB raw, 5.9TB clean dataset
- Look for it in /mnt/oarc-pool4/ZSK-20161001/ at an analysis server near you

DITL 2017

- DITL 2017! Thank you contributors!
  - 9.5TB raw, 5.3TB clean, 12.8TB total
  - And counting….more data is still coming and to be processed
  - Probably the largest ever by volume
  - We also happened to collect data on a 40Mb/s ODVR DDOS….thanks! ☺
- First demanding use of fs4’s CPUs for analysis of incoming data and it has responded well
- Look for it in /mnt/oarc-pool4/DITL-20170411/ at an analysis server near you

DITL 2017 Notes

- Many notes for researchers on DITL 2017 web page than in the previous DITL collections
- This was the first DITL to optionally offer uploading to a site outside the USA
  - No takers
  - We won’t be continuing this in the future
- Coverage maps are being revised due to delayed uploading

General DITL Notes

- Each DITL collection has a dedicated web page with explanatory notes
- The queries and volumes table presented for each collection event have been found to be incorrect due to a SQL bug and are being regenerated
- We have also found that there is additional data from pre-2014 not represented in CLEAN/ sub trees so these are likely to get a fresh reprocessing
  - Yet no one noticed….
- Level3 DNS contributions in years 2010, 2011 and 2012 are not pcap-format, merely BIND logs and are not counted as part of the queries and volumes table.

Network effects of DITL 2017

- There was some concern expressed over whether we could handle network demands
- Past IPerf3 testing had shown line-rate support of 1Gb/s of traffic, so no sweat

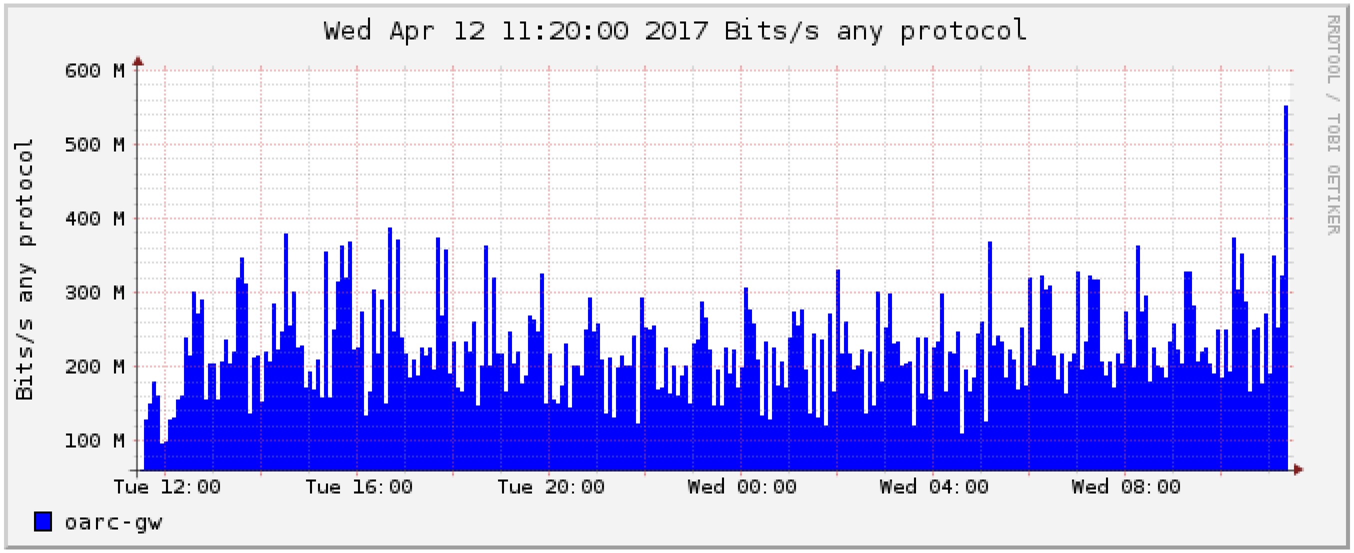

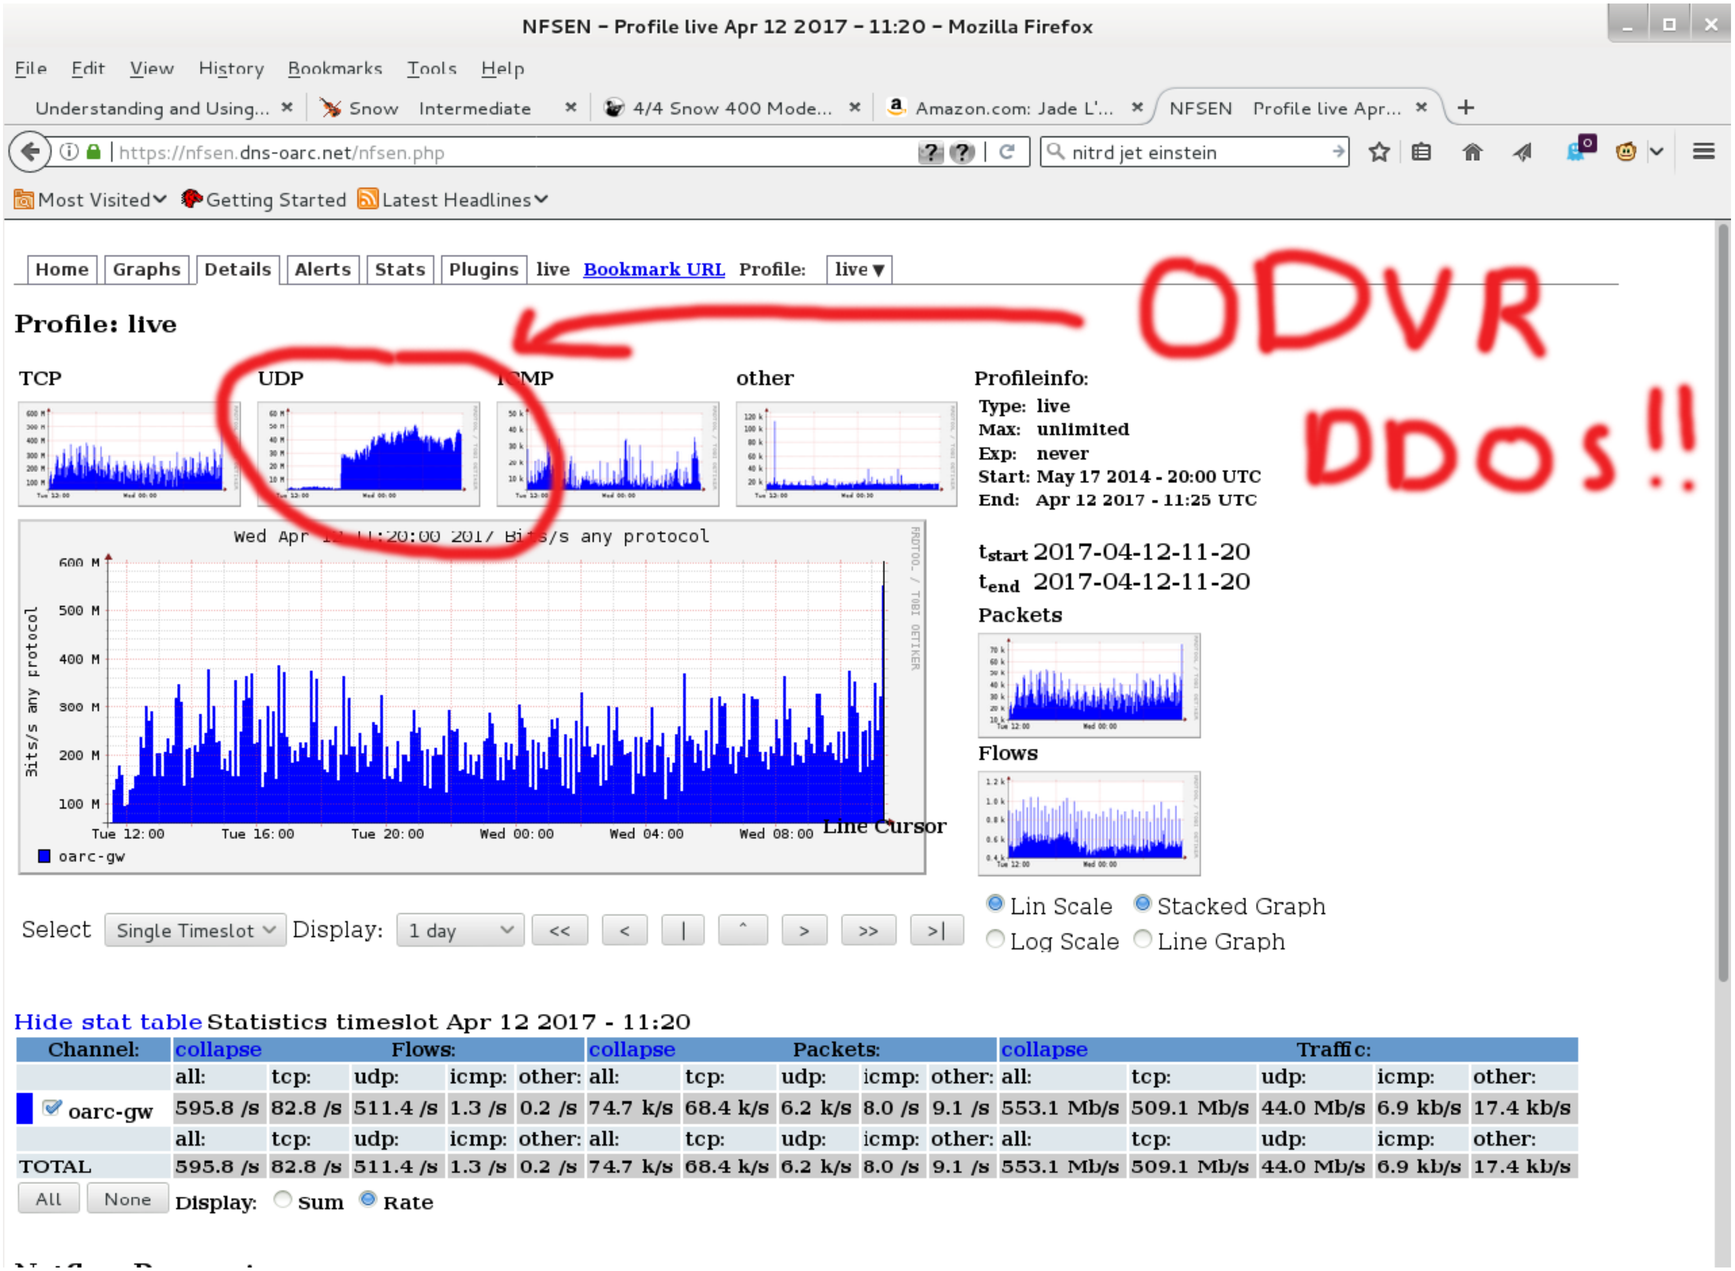

Profile: live

Profileinfo:
Type: live
Max: unlimited
Exp: never
Start: May 17 2014 - 20:00 UTC
End: Apr 12 2017 - 11:25 UTC

t_{start} 2017-04-12-11-20
t_{end} 2017-04-12-11-20

● Lin Scale ● Stacked Graph
○ Log Scale ○ Line Graph

Hide stat table Statistics timeslot Apr 12 2017 - 11:20

Channel:	collapse	Flows:				collapse	Packets:				collapse	Traffic:			
	all:	tcp:	udp:	icmp:	other:	all:	tcp:	udp:	icmp:	other:	all:	tcp:	udp:	icmp:	other:
☑ oarc-gw	595.8 /s	82.8 /s	511.4 /s	1.3 /s	0.2 /s	74.7 k/s	68.4 k/s	6.2 k/s	8.0 /s	9.1 /s	553.1 Mb/s	509.1 Mb/s	44.0 Mb/s	6.9 kb/s	17.4 kb/s
	all:	tcp:	udp:	icmp:	other:	all:	tcp:	udp:	icmp:	other:	all:	tcp:	udp:	icmp:	other:
TOTAL	595.8 /s	82.8 /s	511.4 /s	1.3 /s	0.2 /s	74.7 k/s	68.4 k/s	6.2 k/s	8.0 /s	9.1 /s	553.1 Mb/s	509.1 Mb/s	44.0 Mb/s	6.9 kb/s	17.4 kb/s

All None Display: ○ Sum ● Rate

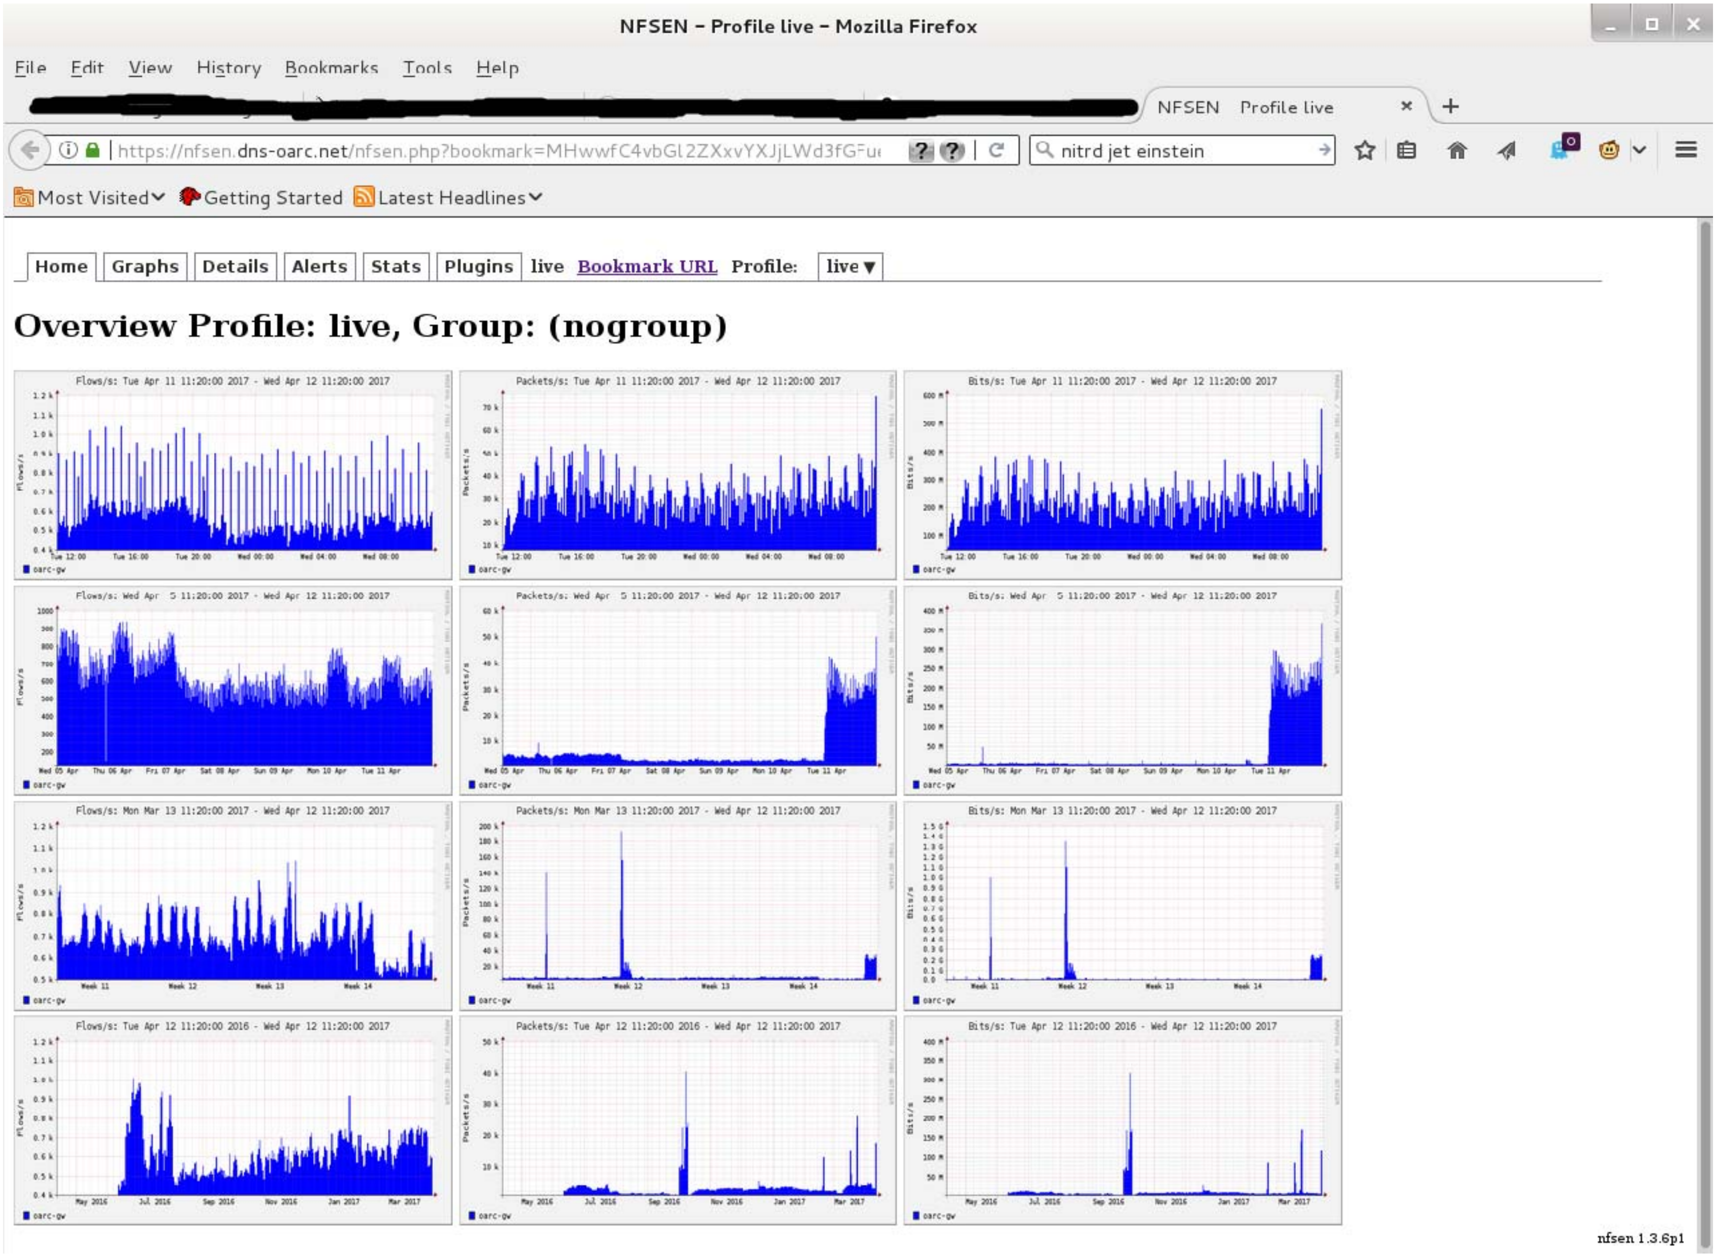

NFSEN – Profile live – Mozilla Firefox
File
Edit
View
History
Bookmarks
Tools
Help
[redacted]
NFSEN Profile live
https://nfsen.dns-oarc.net/nfsen.php?bookmark=MHwwfC4vbGl2ZXxvYXJjLWd3fG=u
nitrd jet einstein
Most Visited
Getting Started
Latest Headlines
Home
Graphs
Details
Alerts
Stats
Plugins
live
Bookmark URL
Profile:
live
Flows/s: Tue Apr 11 11:20:00 2017 - Wed Apr 12 11:20:00 2017
Packets/s: Tue Apr 11 11:20:00 2017 - Wed Apr 12 11:20:00 2017
Bits/s: Tue Apr 11 11:20:00 2017 - Wed Apr 12 11:20:00 2017
Flows/s: Wed Apr 5 11:20:00 2017 - Wed Apr 12 11:20:00 2017
Packets/s: Wed Apr 5 11:20:00 2017 - Wed Apr 12 11:20:00 2017
Bits/s: Wed Apr 5 11:20:00 2017 - Wed Apr 12 11:20:00 2017
Flows/s: Mon Mar 13 11:20:00 2017 - Wed Apr 12 11:20:00 2017
Packets/s: Mon Mar 13 11:20:00 2017 - Wed Apr 12 11:20:00 2017
Bits/s: Mon Mar 13 11:20:00 2017 - Wed Apr 12 11:20:00 2017
Flows/s: Tue Apr 12 11:20:00 2016 - Wed Apr 12 11:20:00 2017
Packets/s: Tue Apr 12 11:20:00 2016 - Wed Apr 12 11:20:00 2017
Bits/s: Tue Apr 12 11:20:00 2016 - Wed Apr 12 11:20:00 2017
oarc-gw
nfsen 1.3.6p1

Overview Profile: live, Group: (nogroup)

Other News for Data Archive

- Fs5 was retrofitted to bring online 22TB plus fs4's former 45TB of capacity – nothing has gone to waste – for a total of 67TB capacity
- We are expecting at least two other collections this year related to the KSK, and have capacity beyond that for the DITL collection for 2018
- Fs1 and fs6 mirrors are now full
  - Starting in 2018 data backups will be local again
  - System backups unaffected

Other Notes

- DANE test pages were deployed last year
  - http://dane.dns-oarc.net/
  - Stats will be used to assess utility
  - Projected lifetime of this service was for just one year
- ZFR still going strong
  - We now have 1 zone file each from 1993 and 1994
    - Not in graph due to different data formats
  - Welcome more samples of root zone files missing from the collection predating DNS-OARC
  - Analysis indicates a nearly flat rate of increase for the first time in some years
  - Some gaps are showing in the root zone data, with some days missing – naming is deceptive! Serial numbers need checking
  - /mnt/oarc-pool3/ZFR/ to find them

ZFR

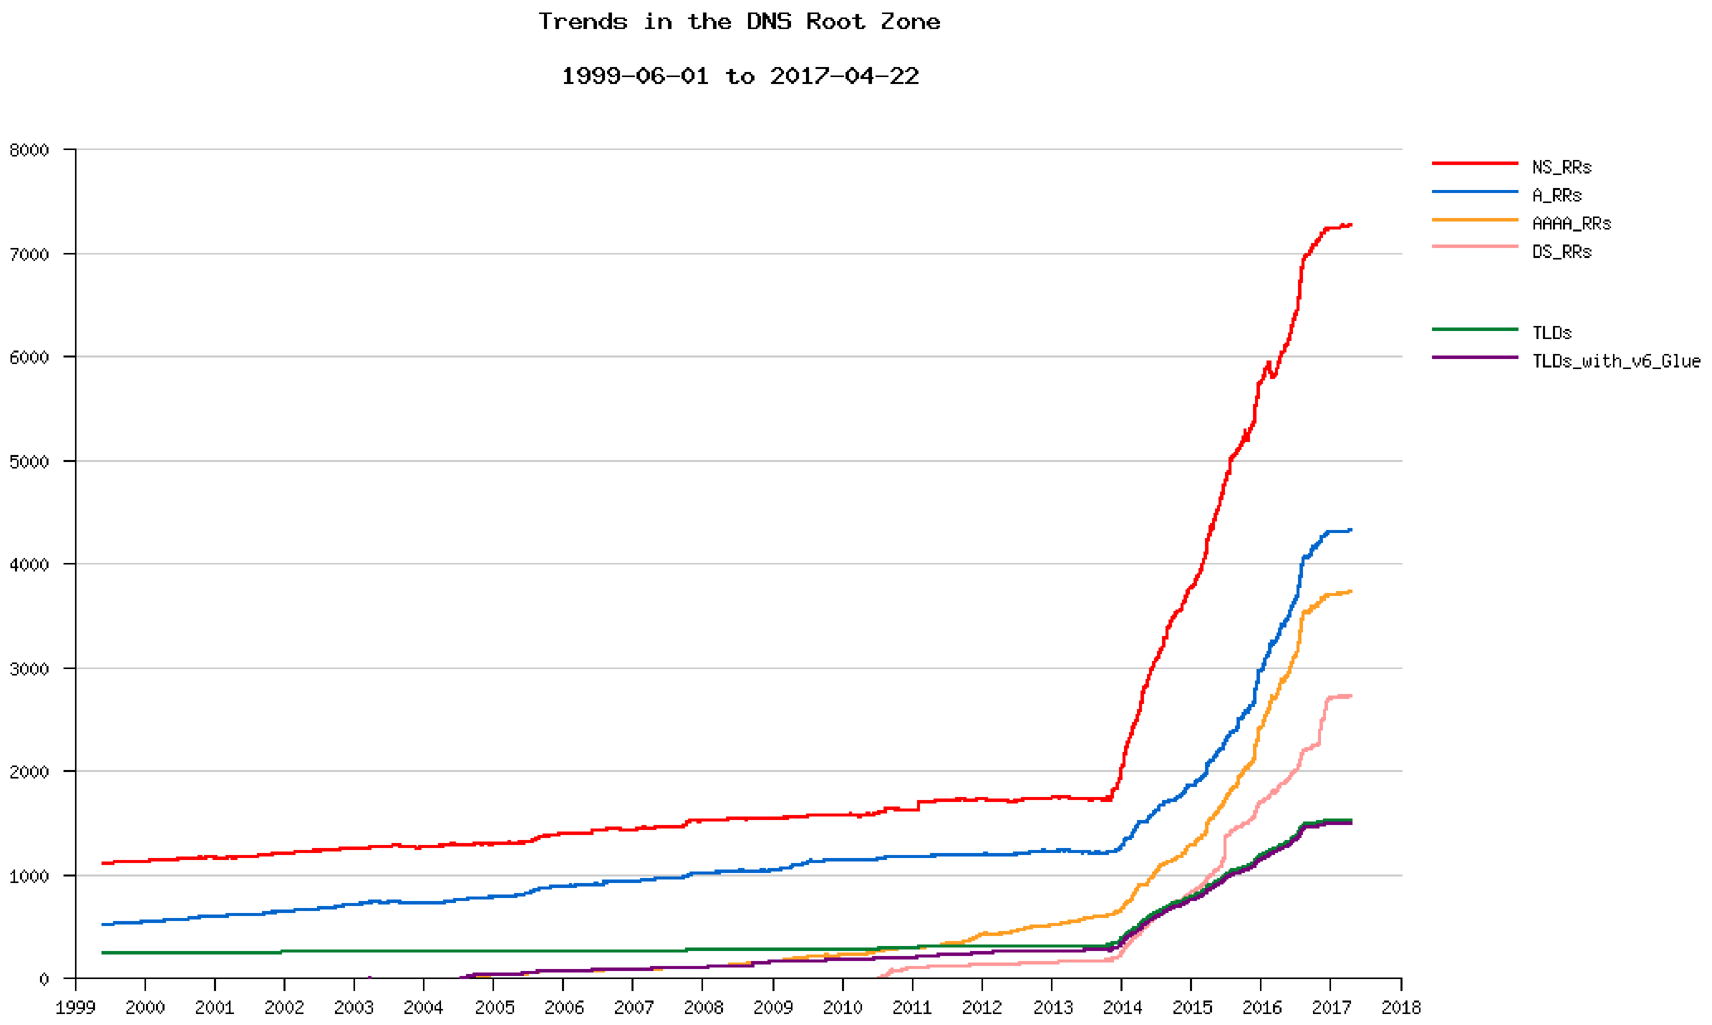
Trends in the DNS Root Zone
1999-06-01 to 2017-04-22
8000
7000
6000
5000
4000
3000
2000
1000
0
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
NS_RRs
A_RRs
AAAA_RRs
DS_RRs
TLDs
TLDs_with_v6_Glue

TLDMon Status

- TLDMon noted 1532 TLDs as of April 24, 2017
- TLDMon has two official sites:
  - https://tldmon.dns-oarc.net/nagios/index.php
  - https://tldmon-se.dns-oarc.net/nagios/index.php
- Experimental TLDMon node at OttIX
  - Monitoring various in-addr.arpa and ip6.arpa zones for fun and scaling, plus new versions of Nagios
  - 1777 zones total being monitored
  - https://tldmon-ca.dns-oarc.net/nagios/index.php
- Perhaps a plugin to detect if a monitored DNS server is actually anycasted

DNS Privacy Servers

- For general info on DNS-over-TLS, see RFC7858
- Over the past summer two such *experimental* servers were turned up at DNS-OARC for protocol features or vulnerabilities testing
  - Using an ISI tool and Unbound both acting as proxies into ODVR
  - Traces and logs available for research and abuse prevention purposes
  - 184.105.193.77 & .78, 2620:ff:c000:0:1::64::25
- See https://www.dns-oarc.net/oarc/services/dnsprivacy
- Not for production use!
- May or may not be up and running always
- Has been abused already…

ODVR

- ODVR maintains a passing acquaintance with its DNSSEC test bed origins
  - It’s really just an open resolver, with all that it entails
  - Looking to increase firewalling protection to limit further abuses
- Operates as a back-end to the TLS-DNS service
- Has complete logging and pcap traces produced daily
- On average the pcaps are usually around 20GB in size daily, so xz is aggressively used to compress them to save space
- DDOS during the DITL 2017 event involved a bunch of spoofed addresses
- /mnt/oarc-pool3/odvr/ houses them

Future

- More capacity growth using 8TB disks or larger in fs1 and fs6 fileservers, or increase the number of fileservers to increase spindle resiliency, perhaps as early as 2018
- Occasional collections if and when they are made will become available.
- Replacement of capture servers using more modern systems and re-integrating the DNS Lab again
- Eventual DR site consolidation
- Potentially offer to install and run AS112 nodes in certain locations.
- Indico major upgrade for late 2018
- Further consideration towards an analysis cluster
- Data catalogue descriptions (ie., what is it?) for each collection
- Re-processing of DITL collections, possibly going back as far as the beginning
- Portal refresh – some call it plastic surgery, others just “restoration.”

“There was another thing…”

- With OS homogeneity between analysis servers it is now possible to make available, if we want to, the following:
  - 160 CPUs
  - 388GB of RAM
  - 10Gb/s Ethernet switch fabric into all analysis servers and file servers on a dedicated network
- In other words, a compute cluster (one location) with an option to form a grid computer (two locations)
- It is possible to harness the 16 DNS Lab systems as well, but requires a lot of additional electrical power
- Plans are to make OpenMPI and/or MPICH libraries available with an eye to deploying Hadoop and Spark as well.
  - Could also look into middleware such as the Globus Toolkit, if it is right for the problem/requirement
- Still a work-in-progress…
- Any takers on naming it?

__END__