DITL & APNIC

George Michaelson

APNIC 26, Christchurch

Overview

- What is DITL?
- How APNIC deployed data-capture
- Outcomes
- Where to from here?

What is DITL

- "**D**ay **I**n **T**he **L**ife" of the Internet
  - 24hours of data capture
  - A snapshot of "what really goes on"
  - Collect, not sample or just measure
- Opportunity for long-life research
  - Compare to past years, future years
  - "what if" questions on data after the event
  - Data archiving, future unplanned uses
    - Rights of use has to be managed

What is DITL (cont)

- Organized by CAIDA
  - Cooperative Association for Internet Data Analysis
    - http://www.caida.org/
  - Safe harbour for data
    - Can observe different T&C to access, publish
    - Long-term data storage
  - Data storage, collation facilitated by OARC
    - https://portal.dns-oarc.net/
    - ‘thumper’ RAID TB filestore
    - S/W support, operations/management

24 hours of data capture!!!

- …actually, it was 48 hours
  - Worldwide coordination required some timezone arithmetic
    - …which some of us got a bit wrong!
      - (ie me…)
  - Previous experience showed getting complete data for a given 24h window was difficult without more data to select from
- Yes. It's a lot of data
  - APNIC's contribution was ~ 300Gb alone

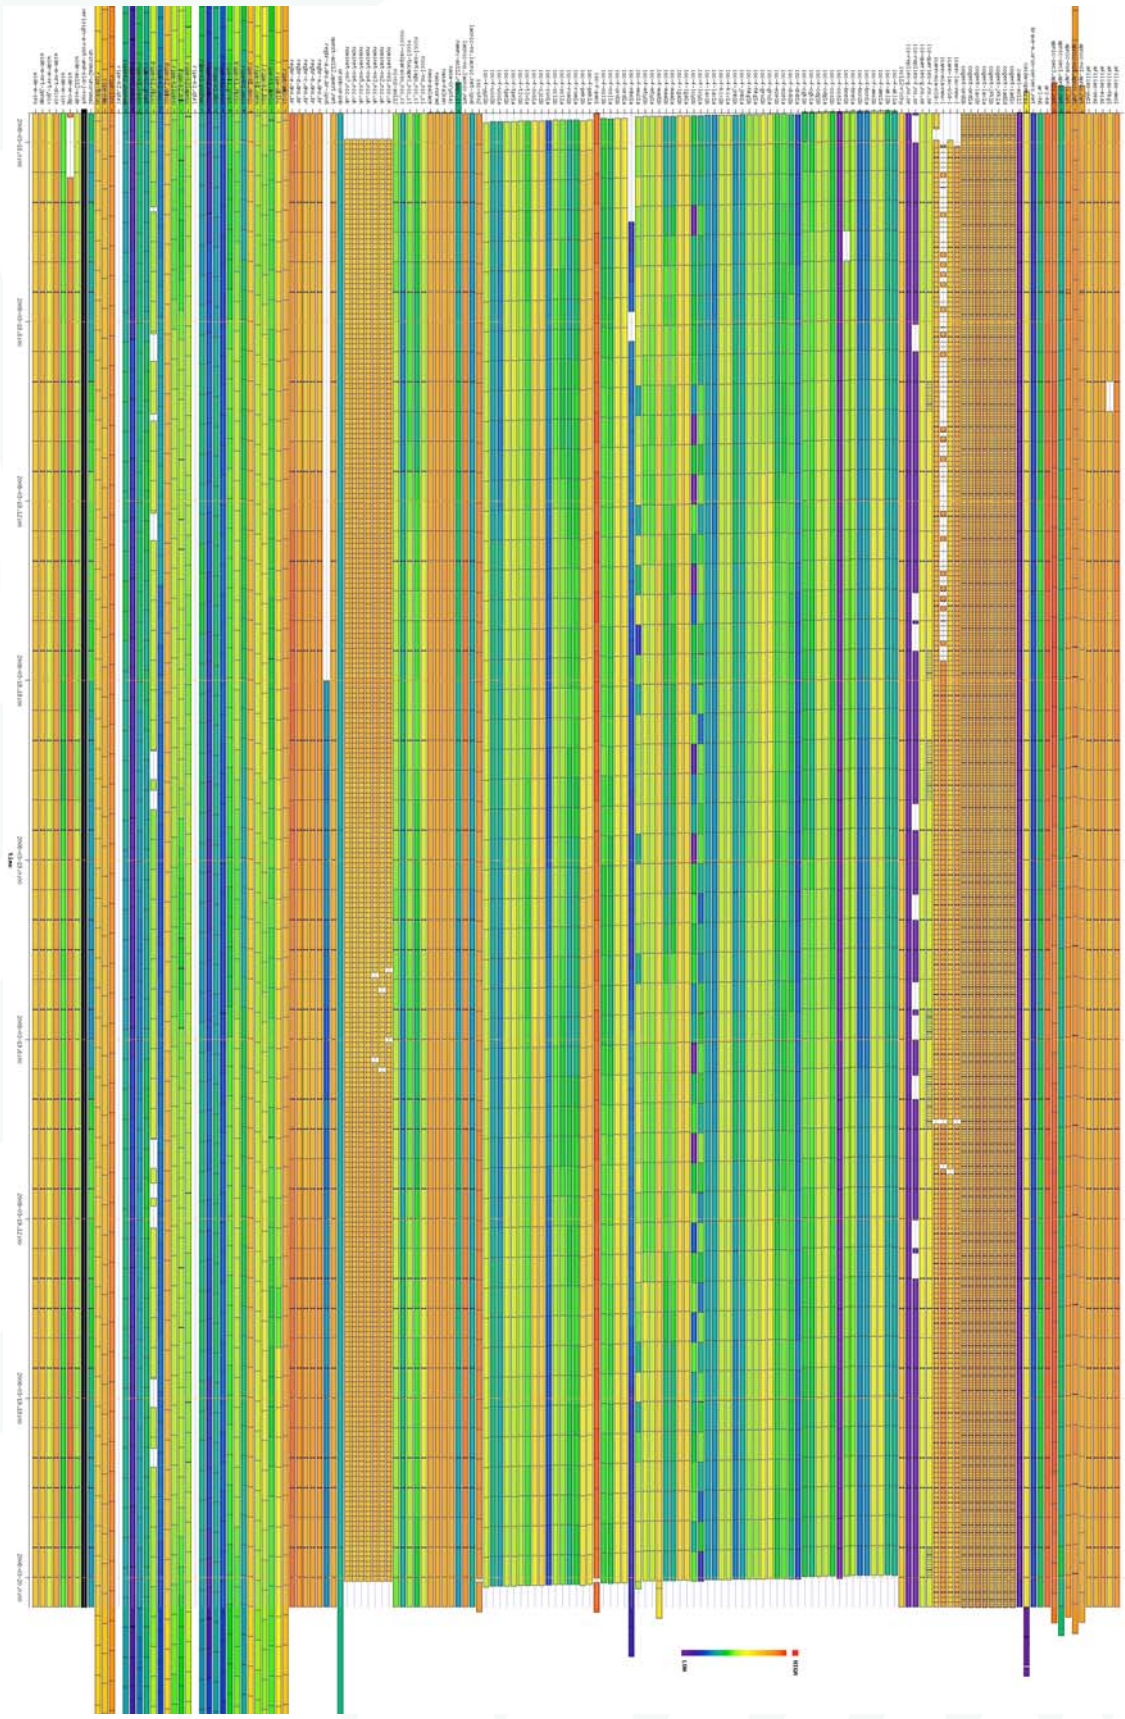

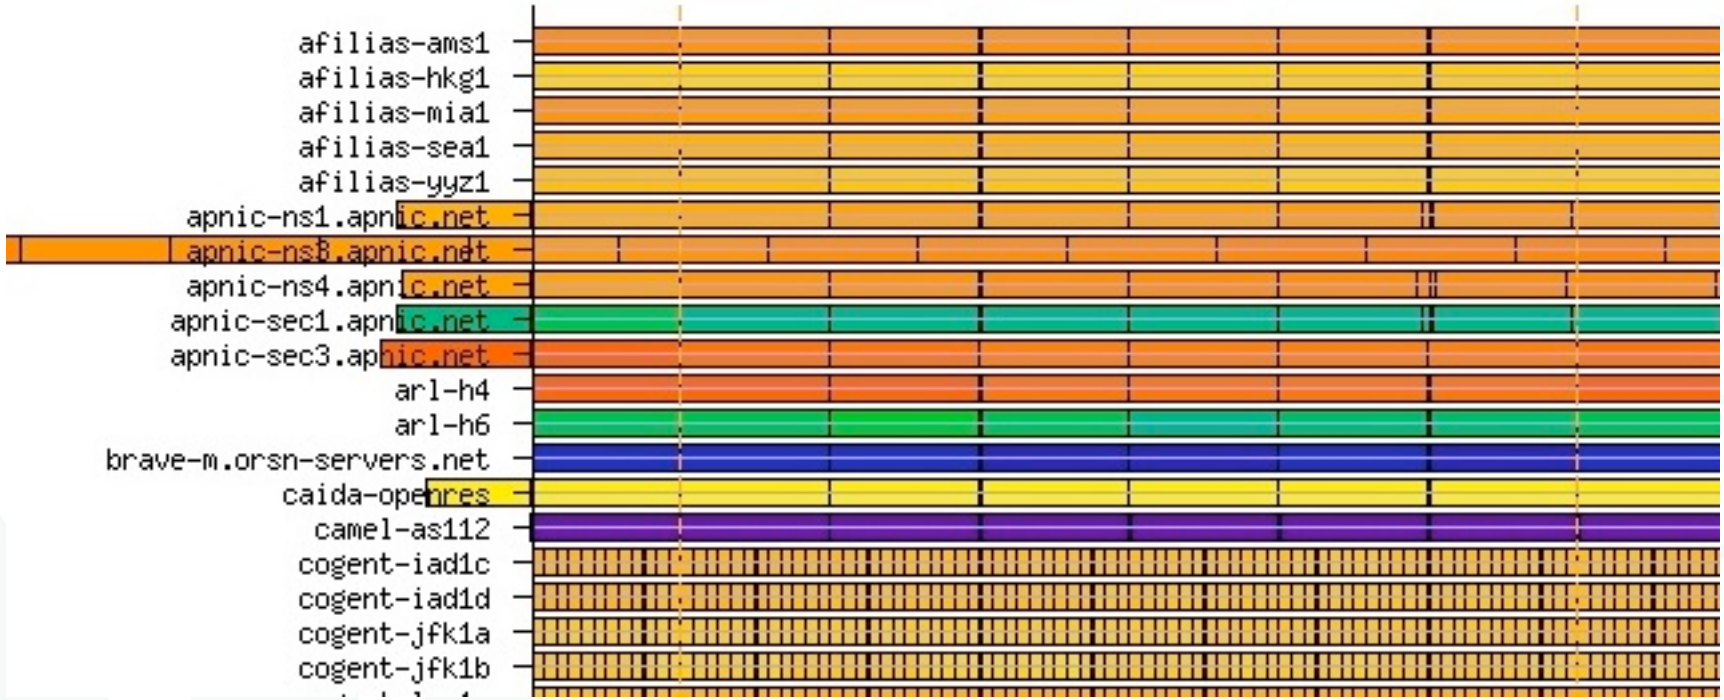
afilias-ams1
afilias-hkg1
afilias-mia1
afilias-sea1
afilias-yyz1
apnic-ns1.apnic.net
apnic-ns3.apnic.net
apnic-ns4.apnic.net
apnic-sec1.apnic.net
apnic-sec3.apnic.net
arl-h4
arl-h6
brave-m.orsn-servers.net
caida-openres
camel-as112
cogent-iad1c
cogent-iad1d
cogent-jfk1a
cogent-jfk1b

Participation

- 156 points of measurement
- 24 agencies
  - Roots, ccTLD, RIR, R&D
- 2Tb of data collected over 48hr+ window
- APNIC's contribution
  - 329Gb
  - All NS, primary (in-region) and secondary (cctLD, other RIR)
  - Passed via SSH to OARC over 96h period

Lessons learned

- Practice makes perfect..
  - Timezone arithmetic
  - Smaller blocks, sent more often, parallelize
- Can’t scale disk to infinity
  - Capture designed for 2+ weeks retention
  - DNS load growth exceeded expectations
- Sample or Measure?
  - Jury still out: both have merits, depending
  - Why not do both?
    - Get samples from capture (used for measurement)

How to capture?

- Originally used on-server tcpdump
  - Higher CPU load, can impact service
  - More chance of packetloss
    - (design of packet filter in kernel drops for the collector, not the real socket endpoint)
  - Servers not designed for this scale of data capture
- Clearly needs re-design for DITL
  - More disk
  - More CPU
  - More flexibility

Passive TAP approach

- Rejected switch-level portspan
  - Imposes switchload, uncomfortable with remote deployment and switch reconfig
- Selected a 1Gig capable TAP, copper
  - Futureproof for re-scaling (current service delivery is 100mbit)
  - Fibre service delivery not yet common in locations APNIC provides service to
  - Vendor-backed 1-packet cost to switchover
    - Dual power supply, failure mode is passive connector
    - Alternate second tap on unit, permits onsite checks
  - No kernel mods, no OS dependencies
    - Libpcap, tcpdump well understood

APNIC DNS measurements

- 5+ year investment in DNS measurements
  - sample-based, not measurement of all DNS
  - 15min duty cycle, 1min packet collector
    - Using tcpdump already
  - Counting by economy, V4/V6, type
- Also OARC “DSC”
  - Measurement of all DNS
  - Distributed collector, central repository model
    - Ssh, rsync, push model to feed repository
  - Libpcap based collector (lib-level tcpdump)
  - XML data retention, mapped into simple CSV file
  - Graphing tools on web (perl)
- Both Ideal to move to TAP based model
  - No loss of historical data series

The APNIC DNS measurement node

- Redhat EL -series OS
- Mirrored 750Gb disks
  - Typically 3-day data retention, up to 10 days for smaller traffic node
- 3-ethernet model (PCI card)
  - Management I/f distinct from capture port
  - Node can measure up to 2 DNS servers onsite
- Packet capture feeds sample, DSC collections
  - Pushed to central collector on management I/f

DITL outcomes

- Published at CAIDA website eg:

  DITL 2008 Analysis
  Sebastian Castro
  Cooperative Association for Internet Data Analysis - CAIDA
  San Diego Supercomputer Center,
  University of California, San Diego

  http://www.caida.org/publications/presentations/2008/oarc_castro_ditlanalysis/oarc_castro_ditlanalysis.pdf

- A Few highlights..

General Stats

		DITL 2007 Root Servers	DITL 2008 Root Servers
Dataset duration		24h	24h
Number of instances		C: 4/4 F: 36/40 K: 15/17 M: 6/6	A: 1/1 C: 4/4 E: 1/1 **(4 nodes)** F: 35/41 H: 2/2 **(v4 and v6)** K: 15/17 L: 2/2 M: 6/6
Query count		3.84 billion	7.56 billion
Unique clients		~2.8 million	~5.6 million
Recursive Queries		17.04 %	11.95 %
TCP	Bytes	1.65%	0.80%
	Packets	2.67%	1.34%
	Queries	~700K	~1.97 million
Queries from RFC1918 address space		4.26%	1.38%
Queries from Bogon address space		0.05%	0.37%

Query rates on IPv6

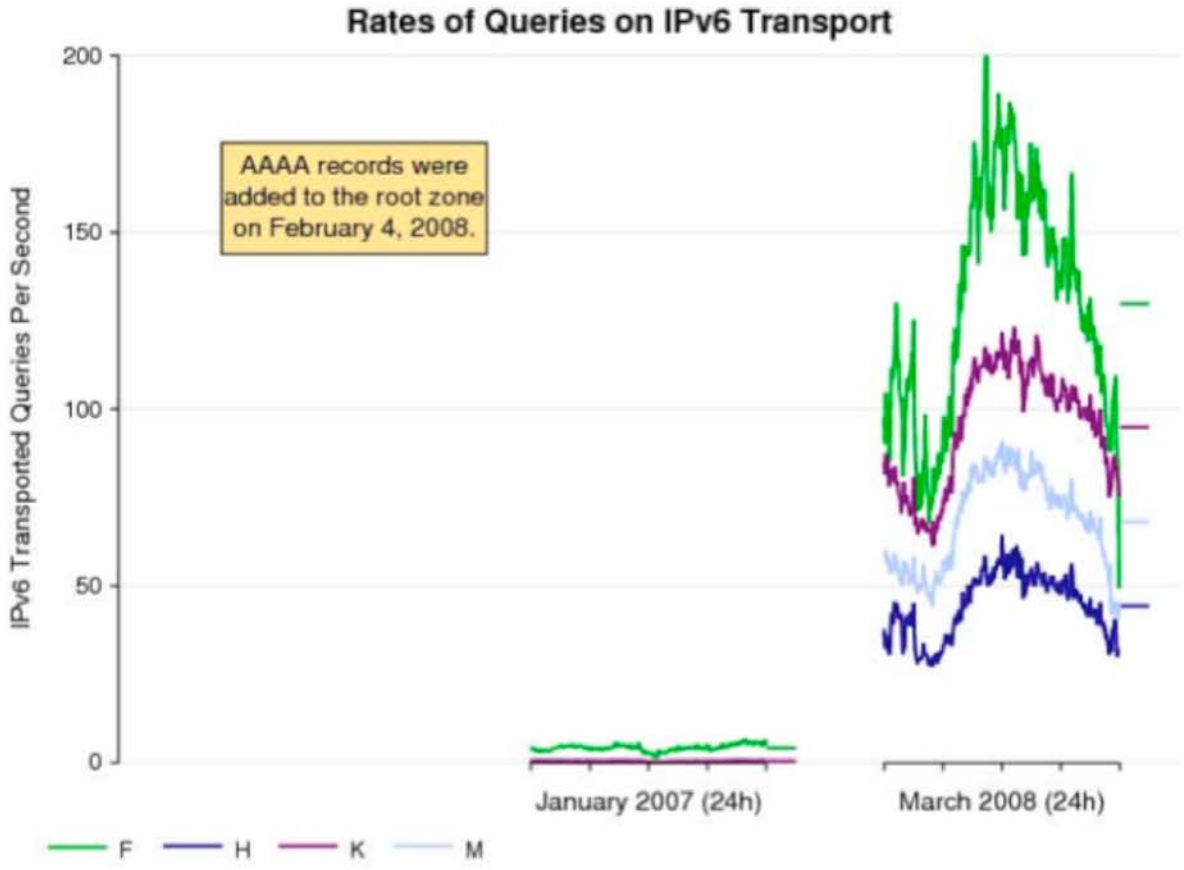

Evolution of EDNS

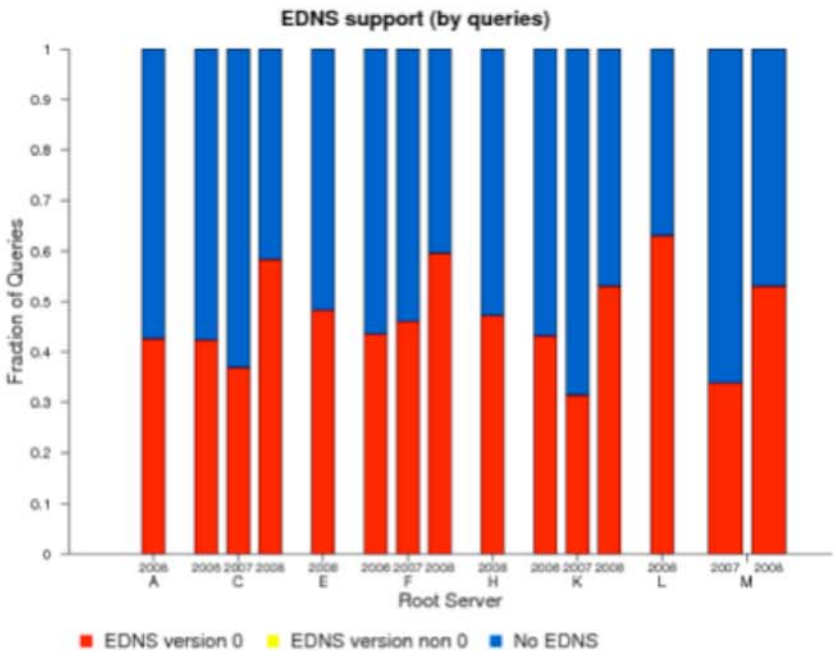

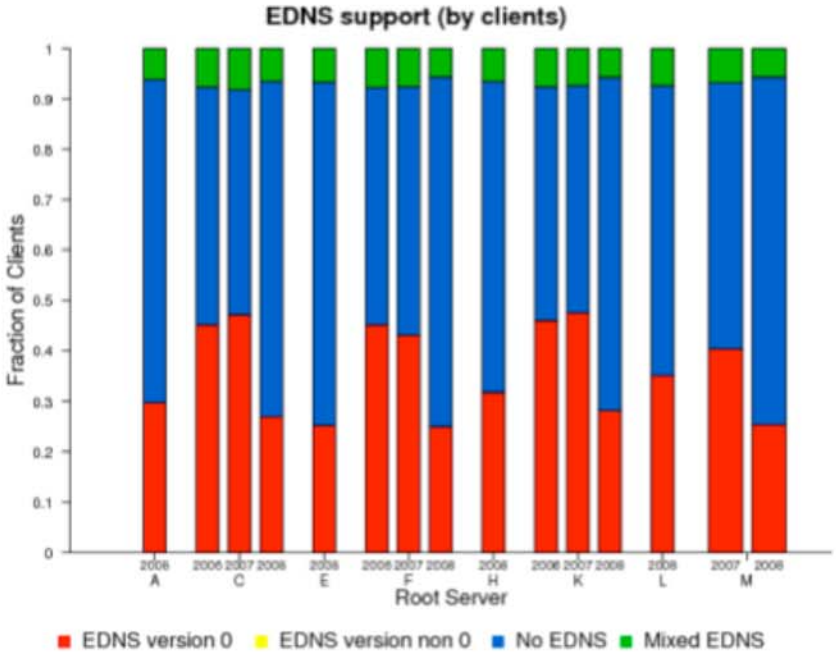

The fraction of queries with EDNS increased in 2008, but the fraction of clients with EDNS support dropped!