DITL 2008 Analysis

Sebastian Castro
secastro@caida.org
CAIDA / NIC Chile

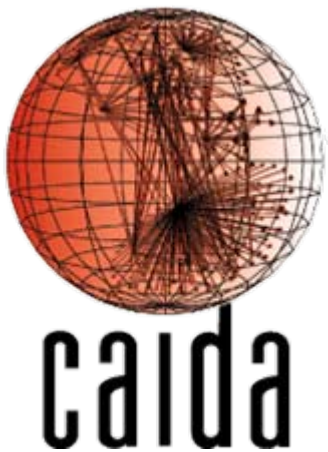

DNS-Ops Workshop. June 4-5, 2008. Brooklyn, NY.

Introduction

- Collection date: Mar 18-19 2008
  - We selected for analysis the traces collected on Mar 19th 2008.
- Participants:
  - Root operators (8 out of 13!)
  - TLD operators
  - RIR
  - All details provided by Duane
- The following analyses are focused only on root server traces
  - The fun part is in Duane's hands ☺

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 Packets Queries	1.65% 2.67% ~700K	0.80% 1.34% ~1.97 million
Queries from RFC1918 address space		4.26%	1.38%
Queries from Bogon address space		0.05%	0.37%

Query rates

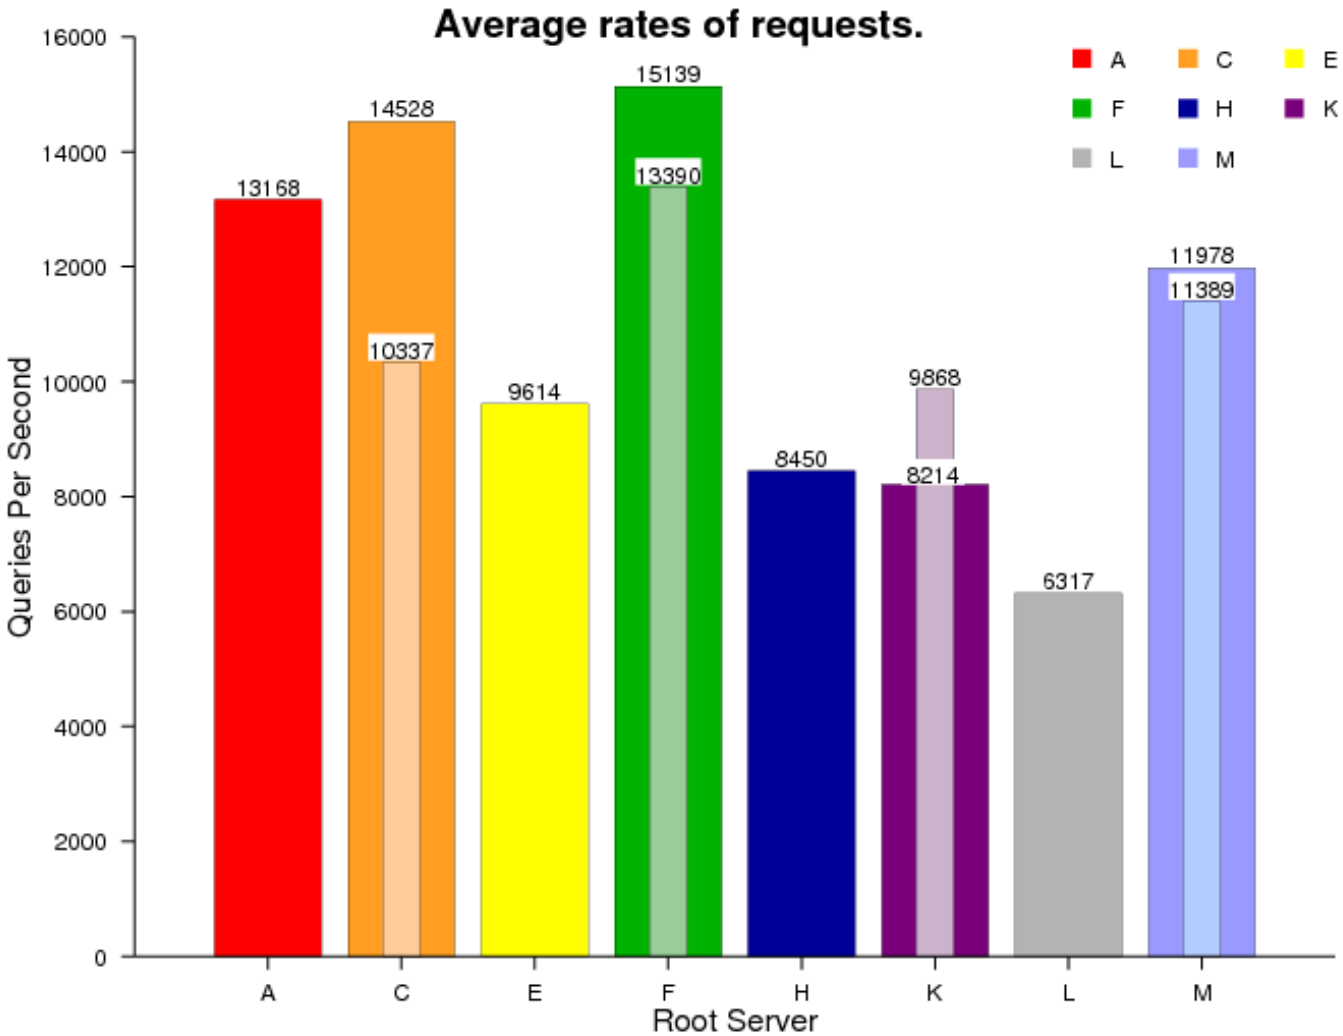

The bar in light color shows the query rates observed in 2007.

Query rates, daily

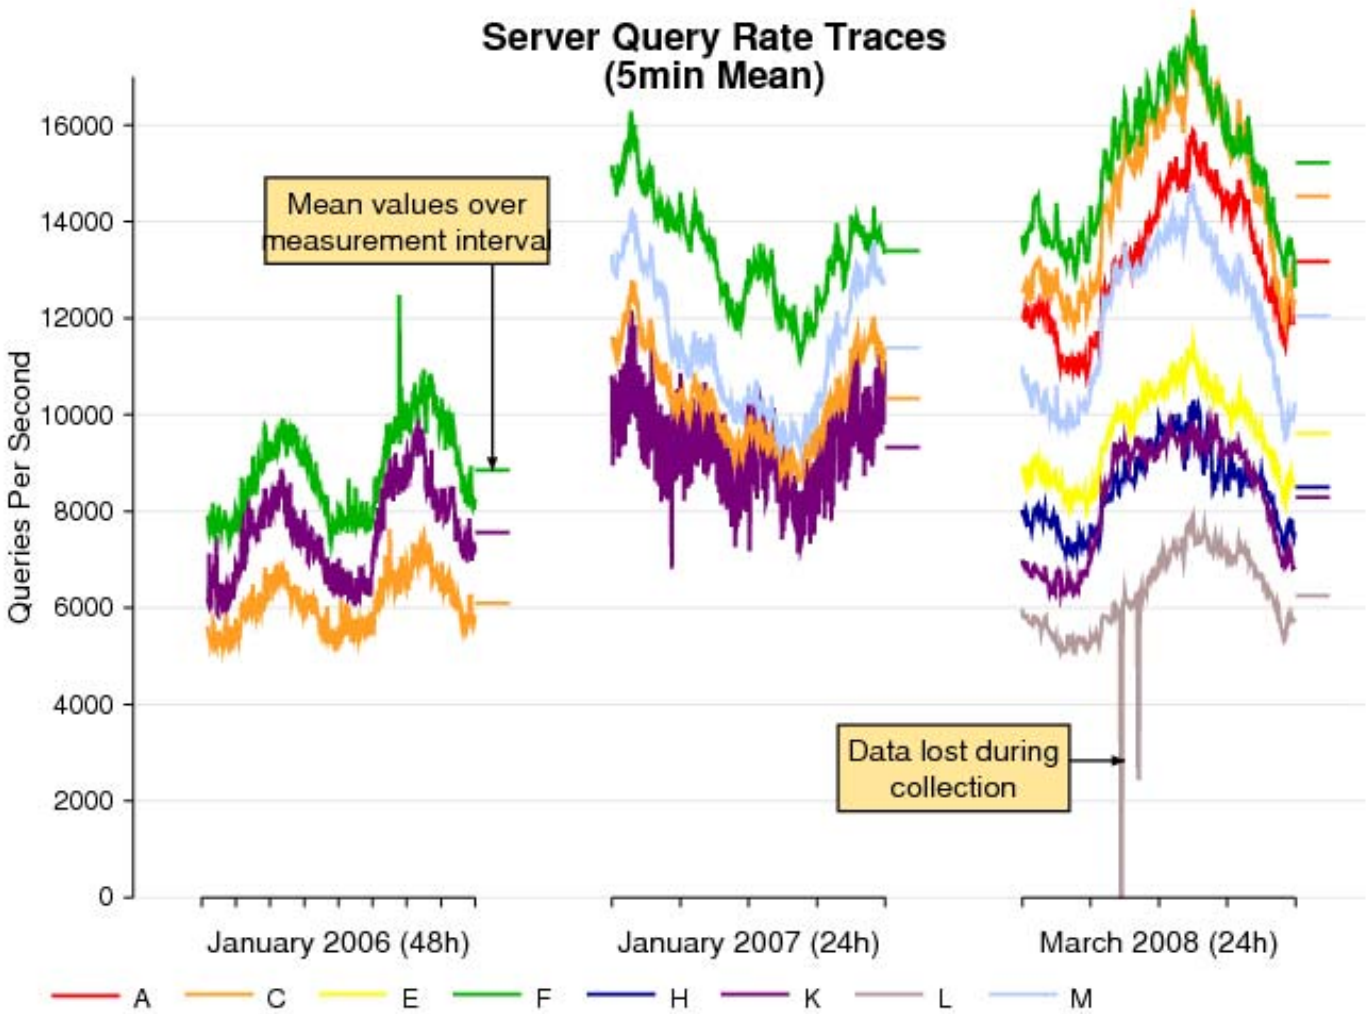

Query rates on IPv6

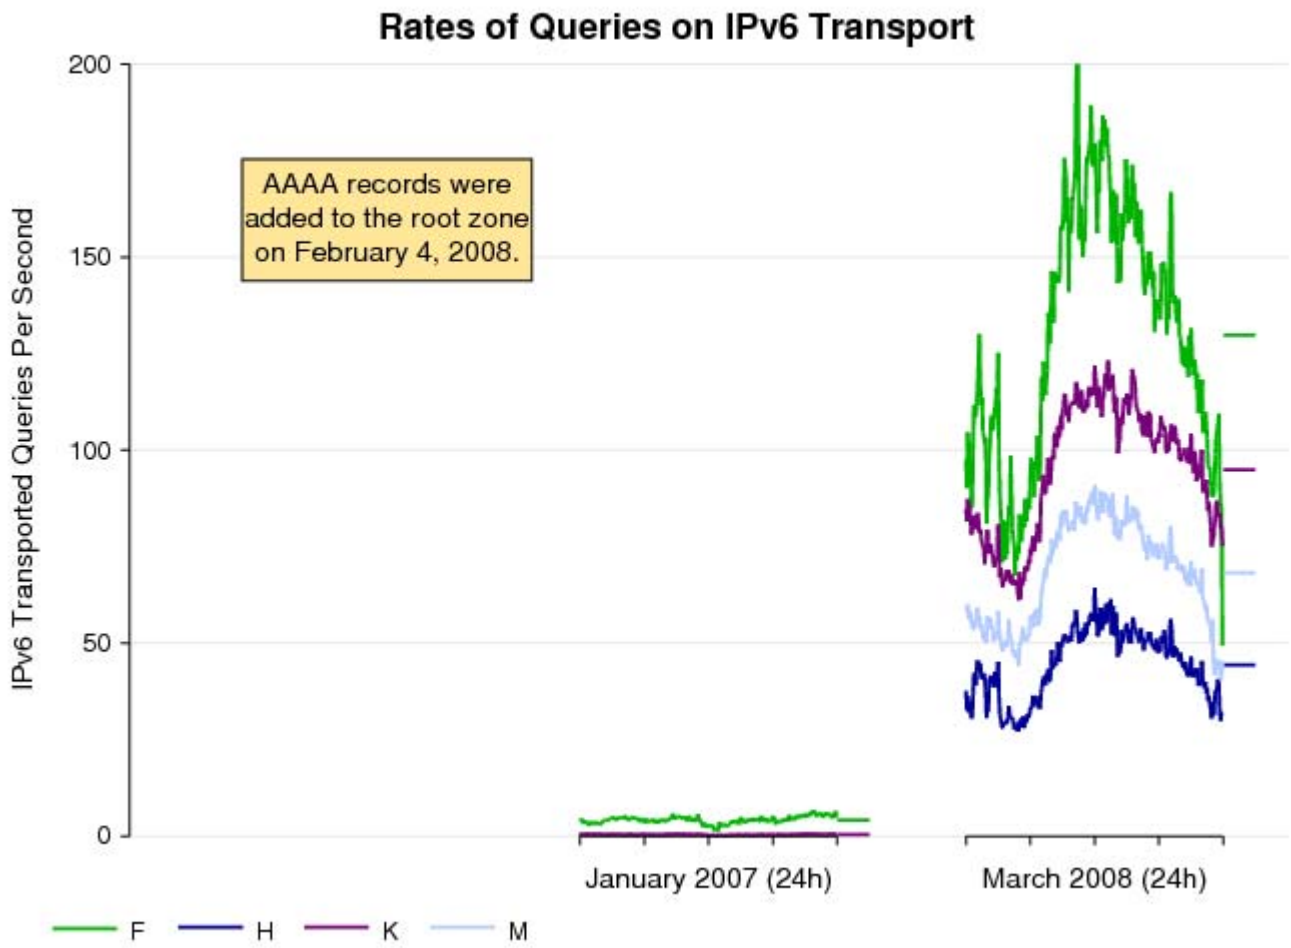

Geography

On the X-axis we have:

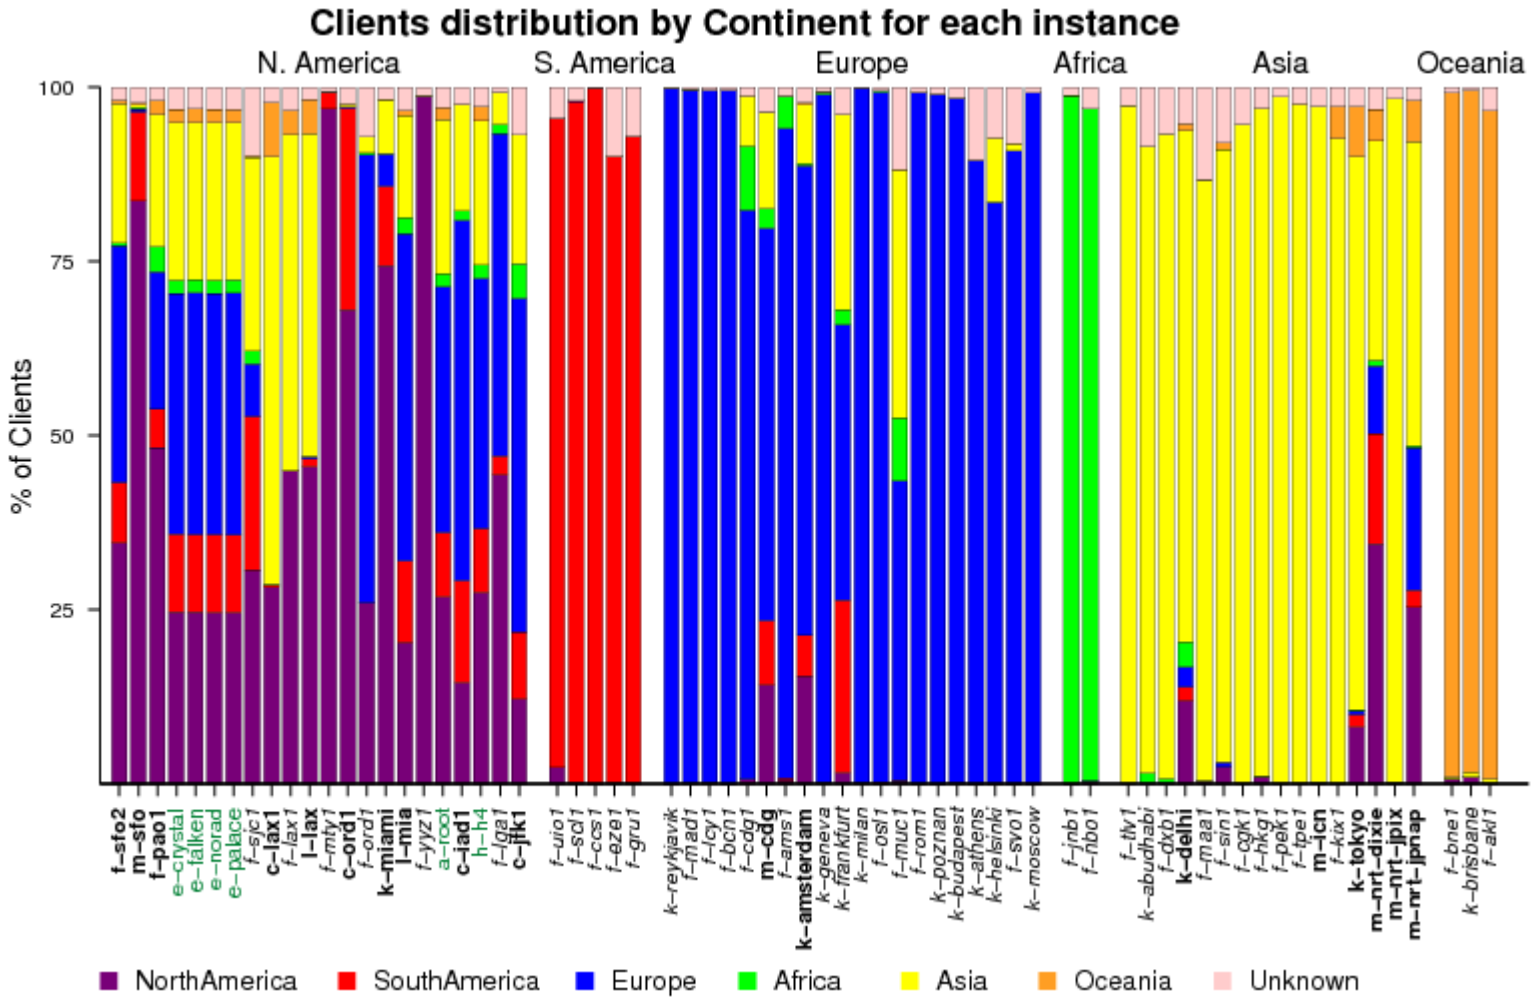

Unicast nodes

Global anycast nodes

Local anycast nodes.

Clients across DNS

- We intersected the list of unique sources querying the root servers with two addresses list
  - Using the traces provided by Afilias (.ORG) and Nominet (.UK)

• 30.01% of the clients were found on all three.

• 49.69% were only seen at the roots.

•Most of them had a query rate < 0.01 qps

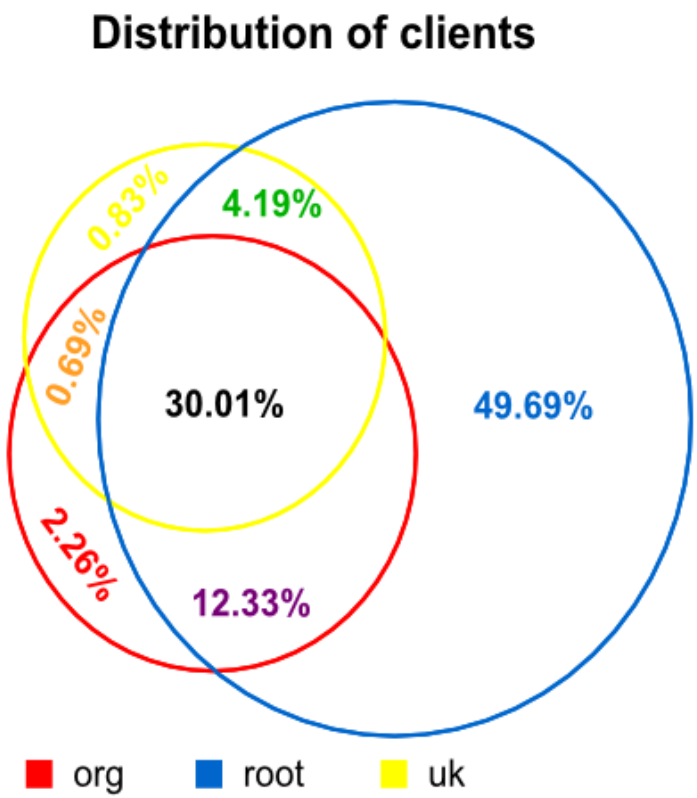

Distribution of queries/clients

The fraction of traffic generated by the heavy hitters (rightmost category) decreased in 2008

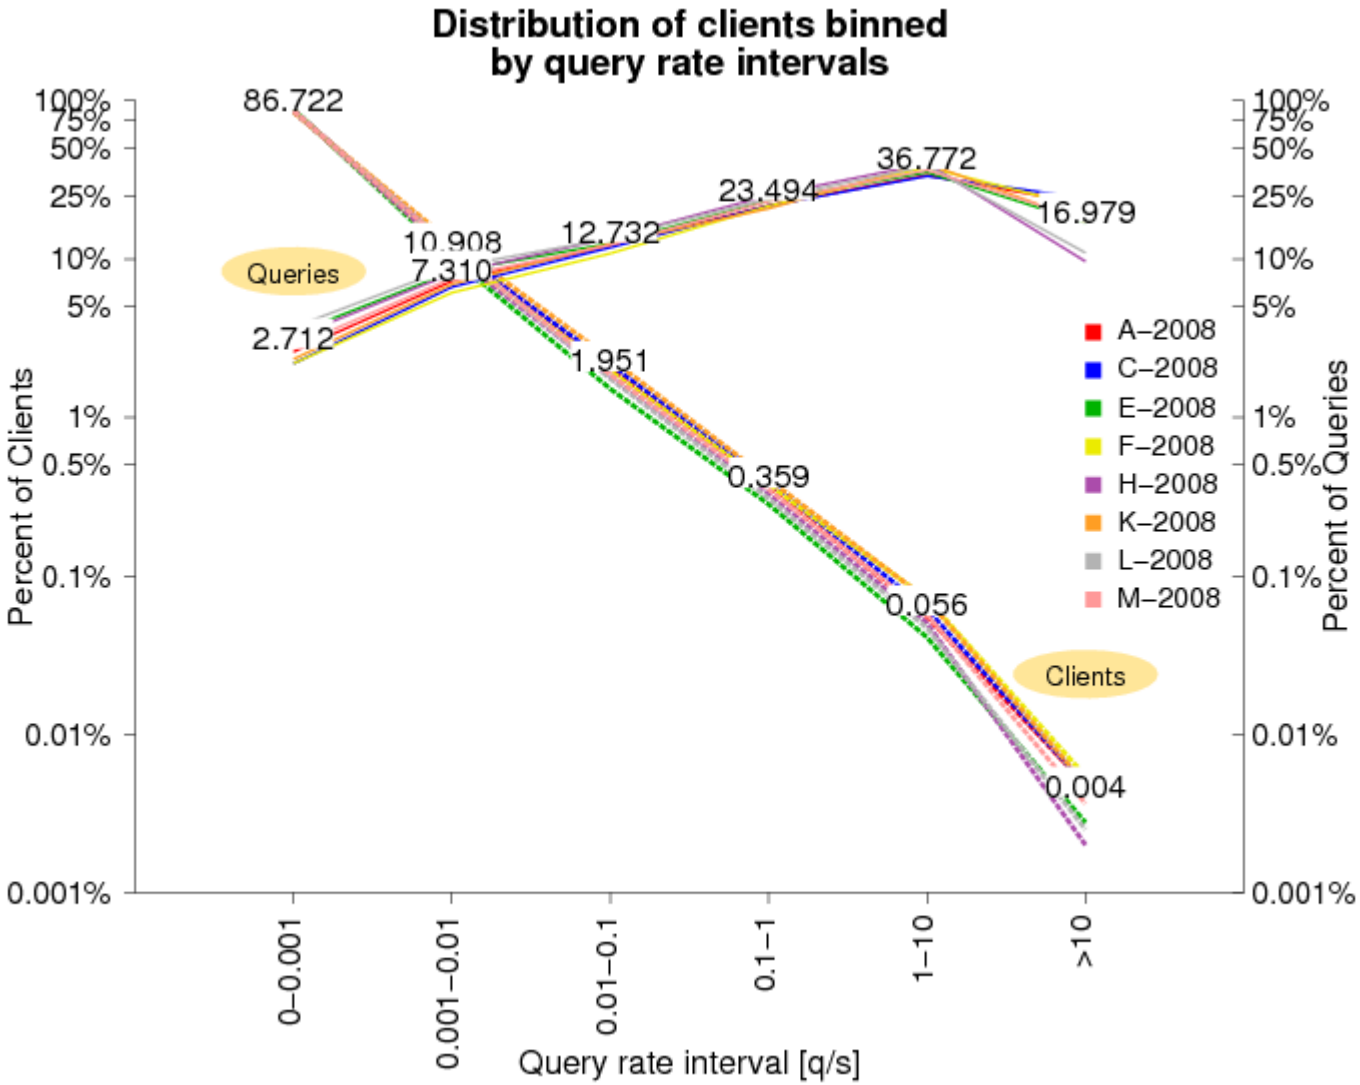

Evolution by query type

We observe an increase in the fraction of AAAA-queries.

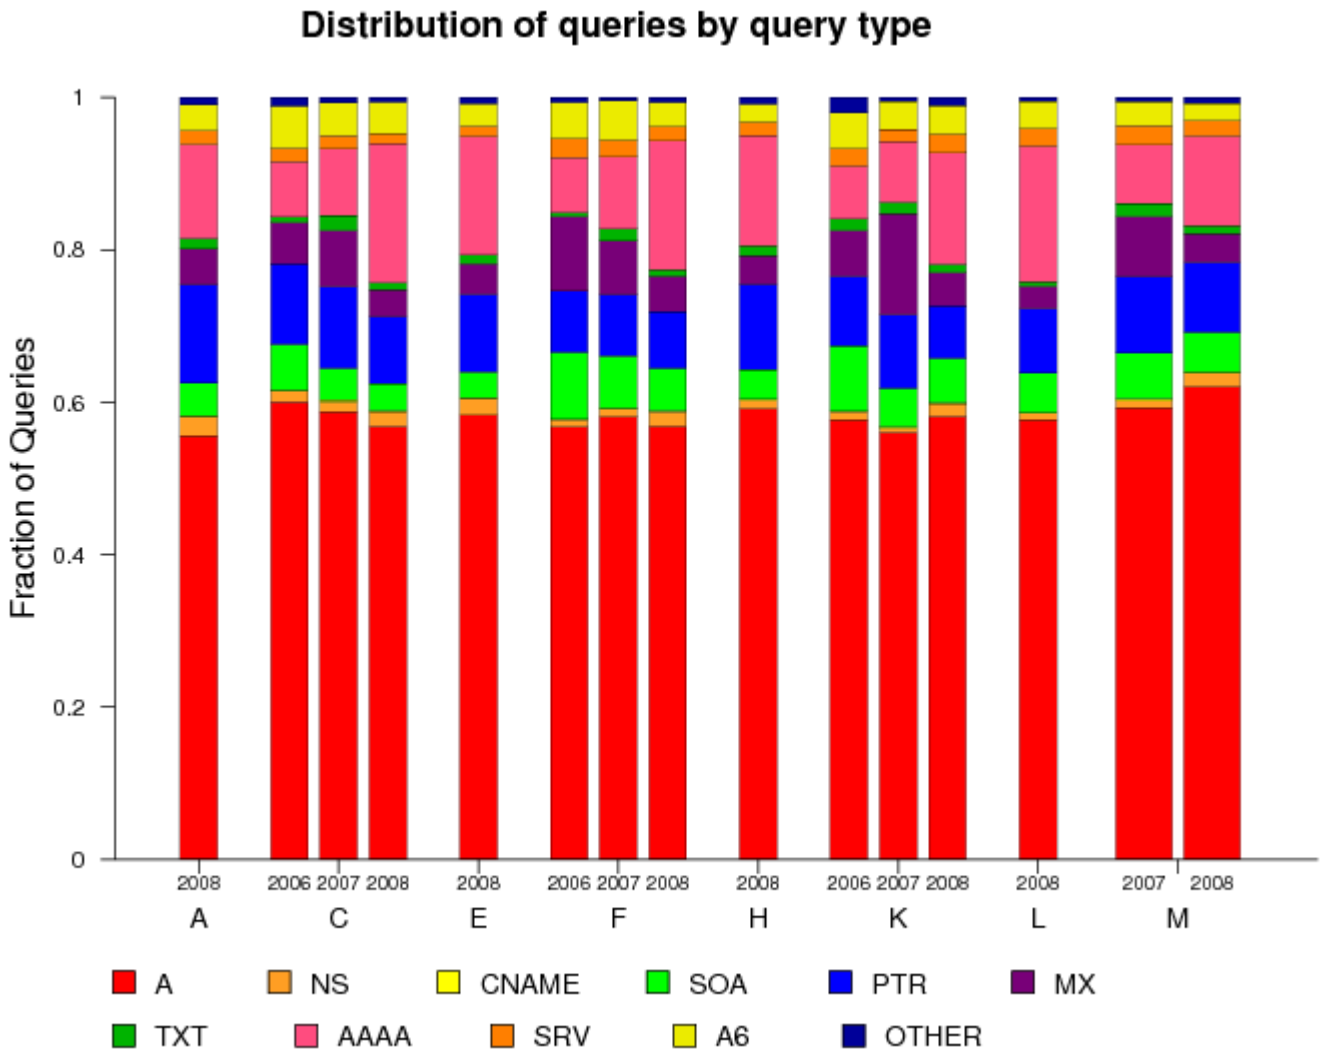

Who is sending the AAAA?

The AAAA queries come mainly from the clients with higher query rates.

5-6% on first two intervals.

10-23% of rightmost interval.

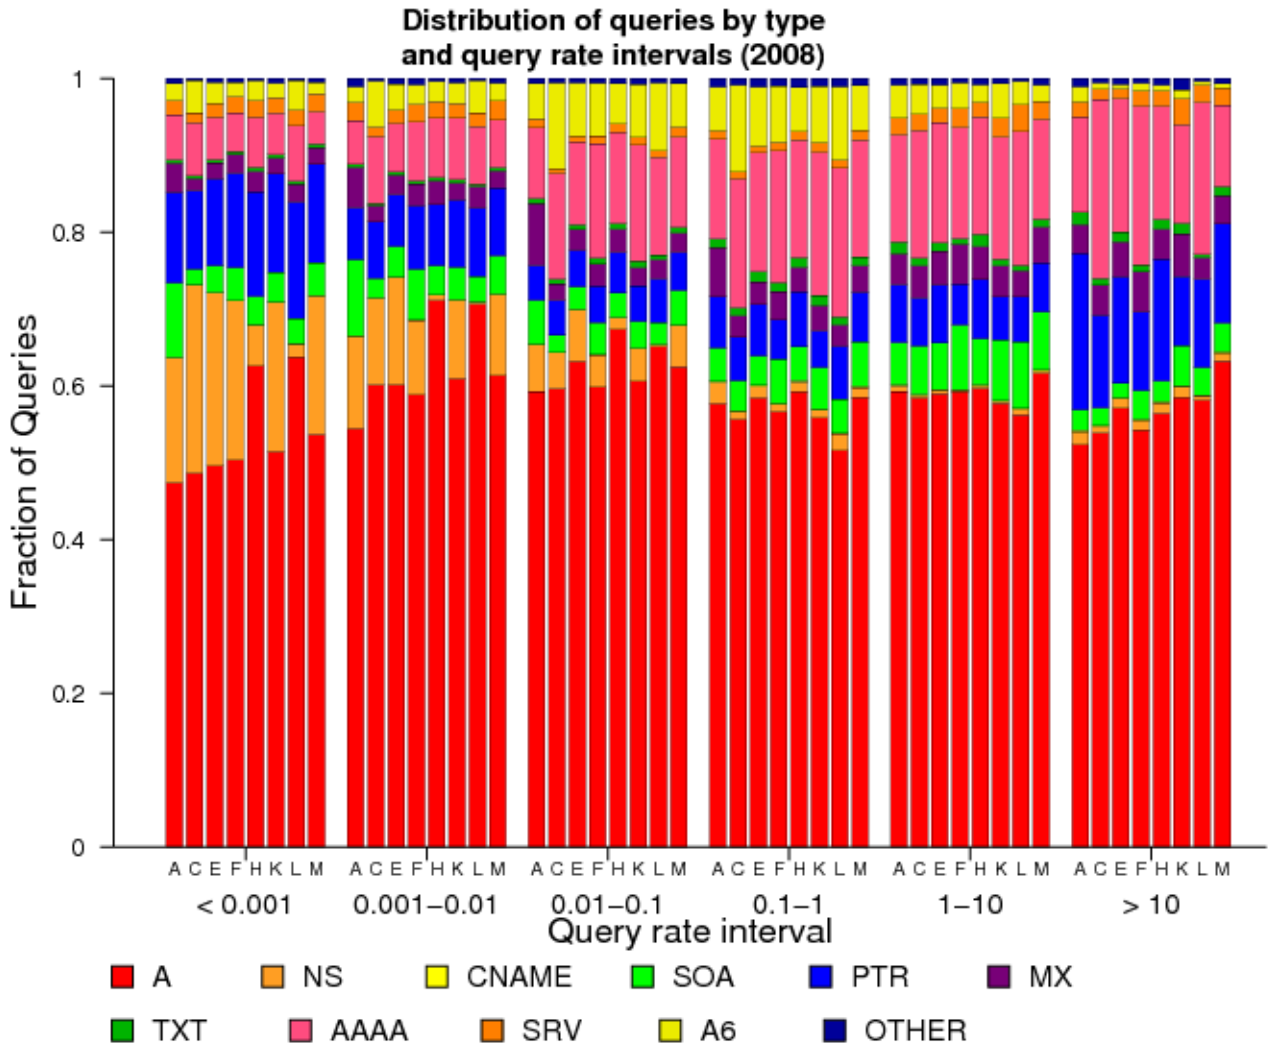

Evolution of EDNS

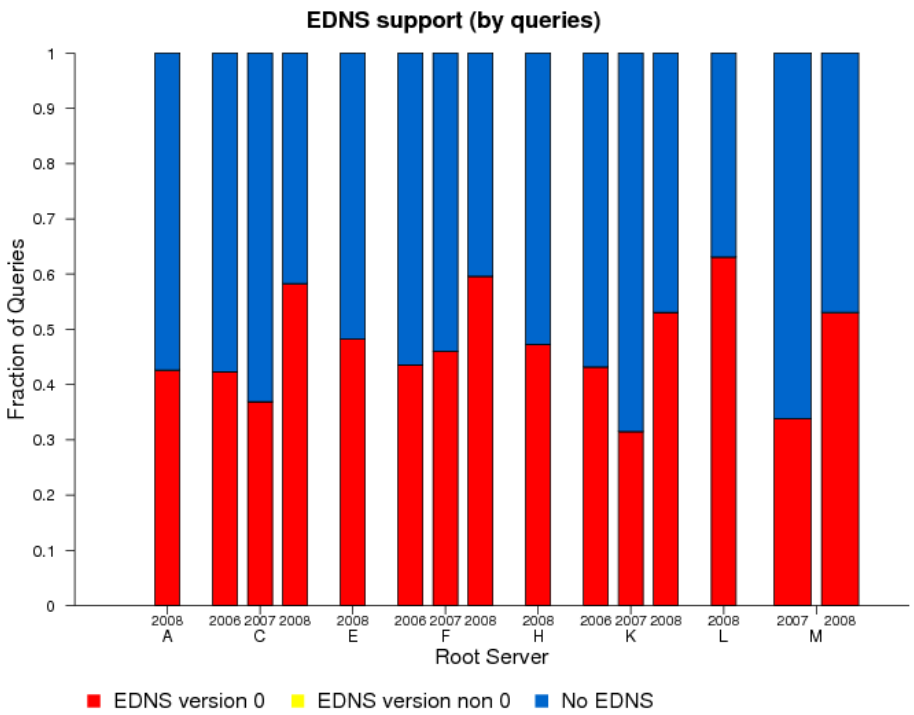

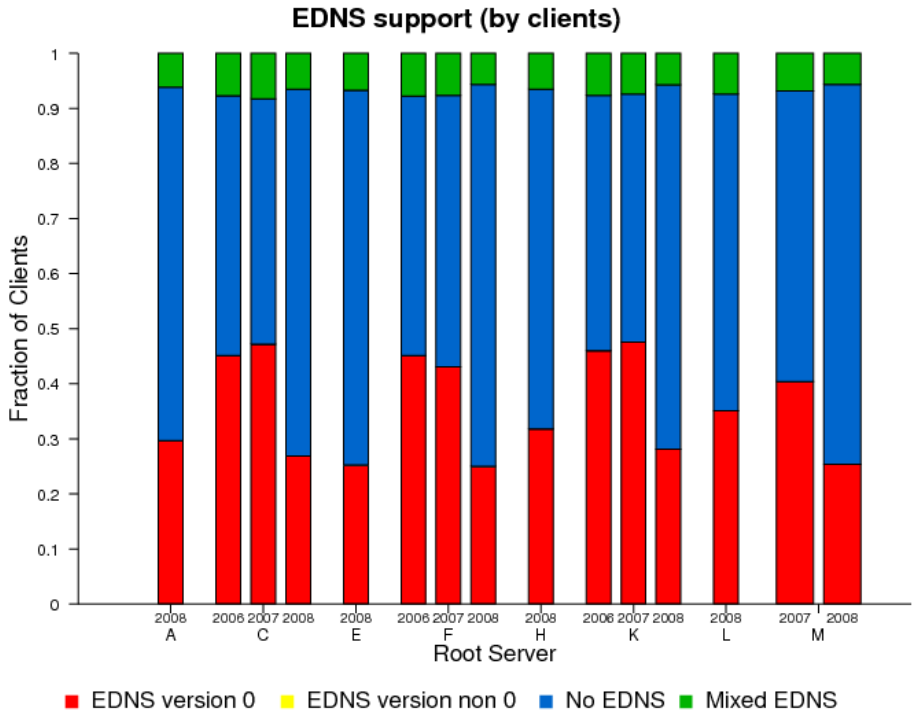

The fraction of queries with EDNS increased in 2008, but the fraction of clients with EDNS support dropped!

Query validity

Category	DITL 2007	DITL 2008	Variation
Unknown Class	0.08	0.09	+0.01
A-for-A	7.02	3.14	-3.88
Invalid TLD	24.73	26.89	+2.16
Non-printable character	0.53	0.05	-0.48
Query name with '_'	0.23	0.15	-0.08
RFC 1918 PTR	0.67	0.47	-0.20

The analysis for 'identical queries', 'repeated queries' and 'referral-not-cached' is not finished yet (has been processing for two weeks)

Taking a closer look: Invalid TLD

TLD	Ranking			Percentage of total queries		
	DITL 2007	DITL 2008	Var.	DITL 2007	DITL 2008	Var.
local	1	1	-	5.018	5.055	+0.037
localhost	2	3	-1	2.205	0.728	-1.477
domain	3	7	-4	0.778	0.550	-0.228
invalid	4	5	-1	0.602	0.629	+0.027
lan	5	4	+1	0.509	0.686	+0.177
belkin	6	2	+4	0.436	0.752	+0.316
home	7	6	+1	0.321	0.594	+0.273
localdomain	8	8	-	0.318	0.336	+0.018
wpad	9	9	-	0.183	0.238	+0.055
txt	10	27	-17	0.182	0.058	-0.124
corp	12	10	+2	0.150	0.233	+0.083

Conclusions

- There is a huge number of clients sending a few queries only to the roots
- There is a increase on the number of AAAA queries.
- The fraction of traffic due to invalid TLD is huge! (25%).
  - Actions toward avoiding those queries reaching the roots could make a real impact.