"Day in the Life 2007" Data Analysis

Sebastian Castro
secastro@caida.org
secastro@nic.cl

CAIDA
NIC Chile

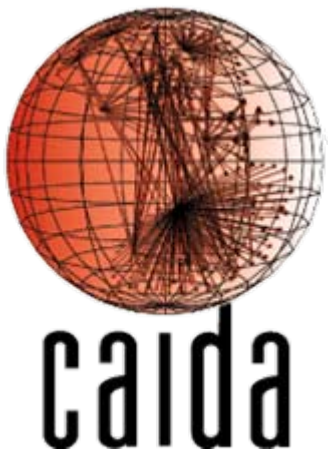

DNS-OARC – November, 2007

DITL 2007

- This report is a continuation of Duane Wessels report on last OARC meeting
- Due to problems on the full DITL 2007 set (~48 hours), we selected the best coverage 24-hour set to work with.
  - January 9th 2007, 12:00:00 to January 10th 2007, 11:59:59
- All further analysis were done using that subset.

Analysis Software

- C++ code
  - Read pcap files, counts and analyze
  - Output SQL and plain text files
  - Range of analysis selected on compilation time
    - Client and query rate, AS/prefix coverage
    - Distribution of queries by query type
    - Node/cloud switching per client
    - Source port analysis
    - EDNS support and EDNS buffer size
    - Invalid queries
    - Others: RD queries, RFC1918/Bogon sources
  - Speed depends on the number of analysis selected
- Data preprocessing: Perl and shell script
- Data plotting using *ploticus*.
- Unexpected feature: Machine crasher
  - Lesson: Don't let a perl geek work on C++ code ☺

Overview

Root	Node	Instances	Avg query rate per instance	Avg query rate all instances	Average client rate
C	Global	4	1723.7 – 3804.3	2584.2	596.2 – 1178.5
F	Global	2	2012.7 – 2563.5	2288.1	589.2 – 1085.1
F	Local	34	3.0 – 1241.1	259.2	1.9 – 344.3
K	Global	5	403.8 – 4864.5	1853.0	134.1 – 1422.2
K	Local	9	3.3 – 259.1	66.9	2.3 – 37.1
M	Global	6	66.8 – 5836.8	1898.1	27.1 – 1930.0

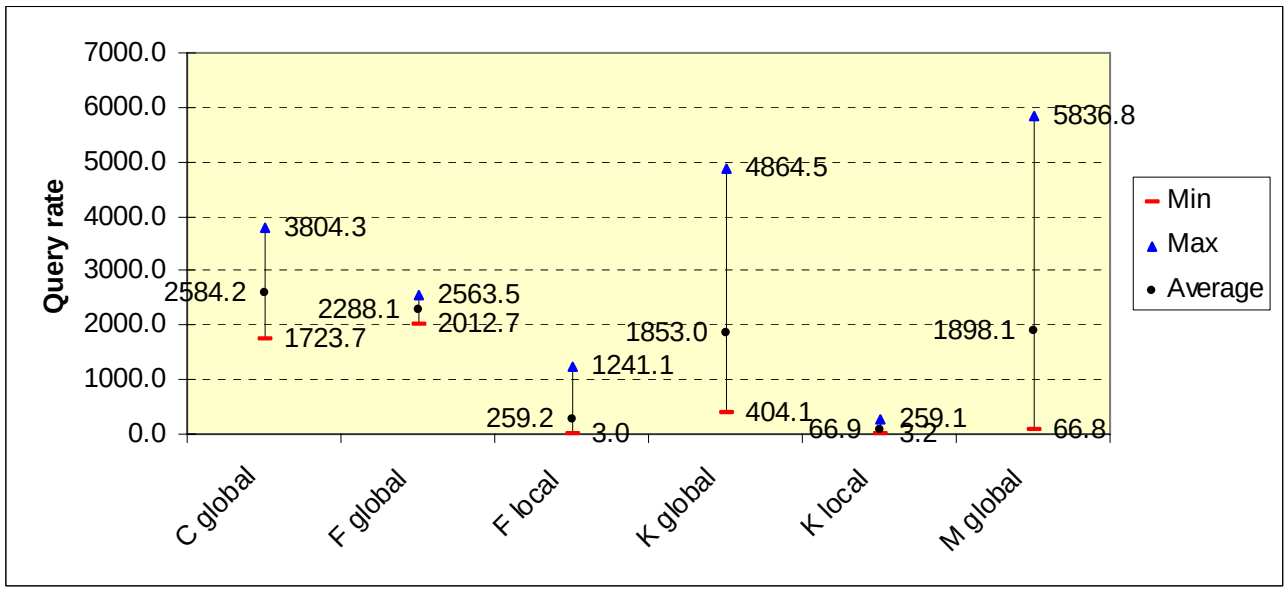

Overview

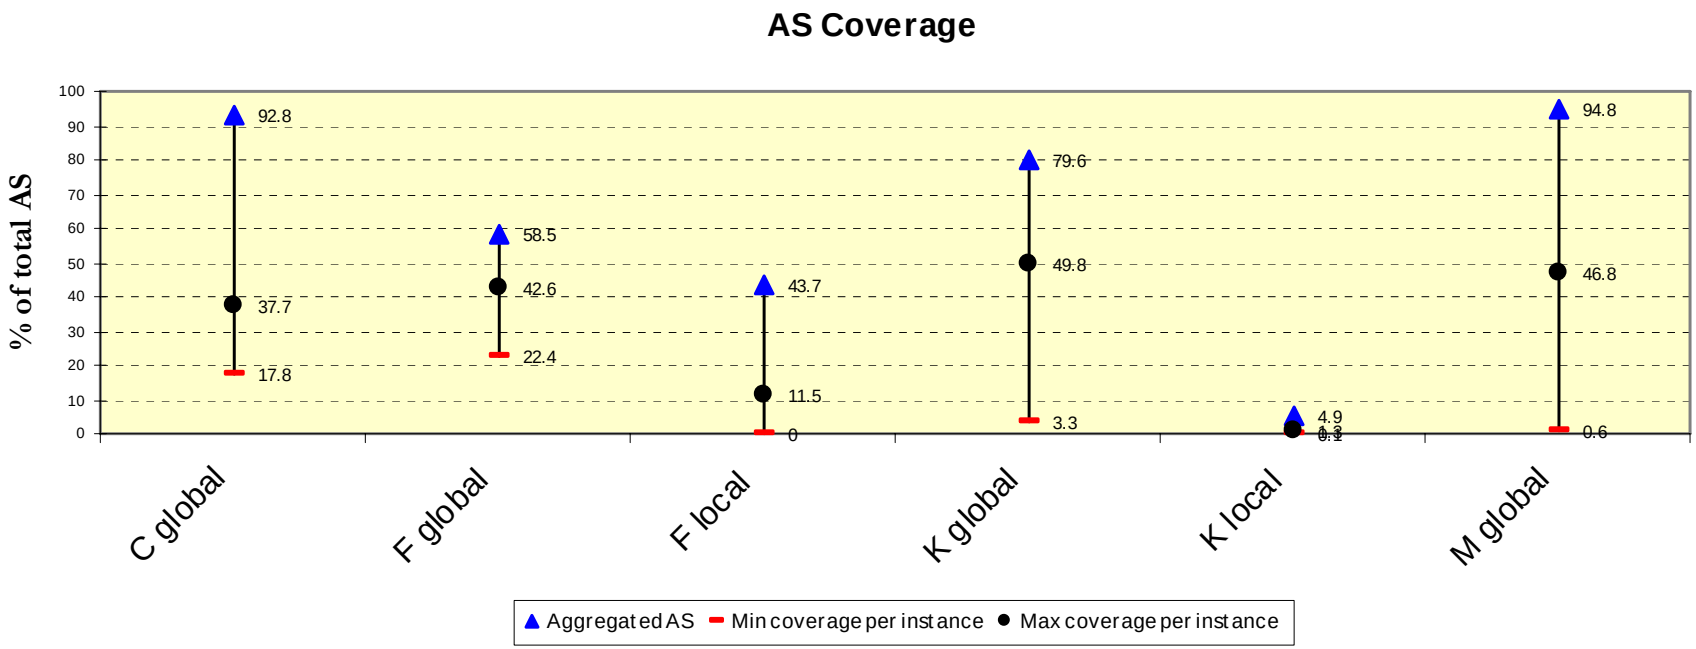
AS Coverage
% of total AS
100
90
80
70
60
50
40
30
20
10
0
92.8
37.7
17.8
58.5
42.6
22.4
43.7
11.5
0
79.6
49.8
3.3
4.9
94.8
46.8
0.6
C global
F global
F local
K global
K local
M global
Aggregated AS
Min coverage per instance
Max coverage per instance

General Stats

	Root 2006 (~48 hours, C, F and K)	Root 2007 (24 hours, C, F, K and M)	ORSN 2007 (24 hours, A and B)
Number of queries	4.92 * 10^9	3.84 * 10^9	4.1 * 10^6
RD traffic (% of queries)	3.61	17.04	11.59
TCP (%)			
Bytes	1.58	1.3	0.17
Packets	2.67	3.2	0.22
Queries	0.0184	0.0064	0.0118
Queries from RFC1918 space	2.15	4.26	0.3

Client Geography

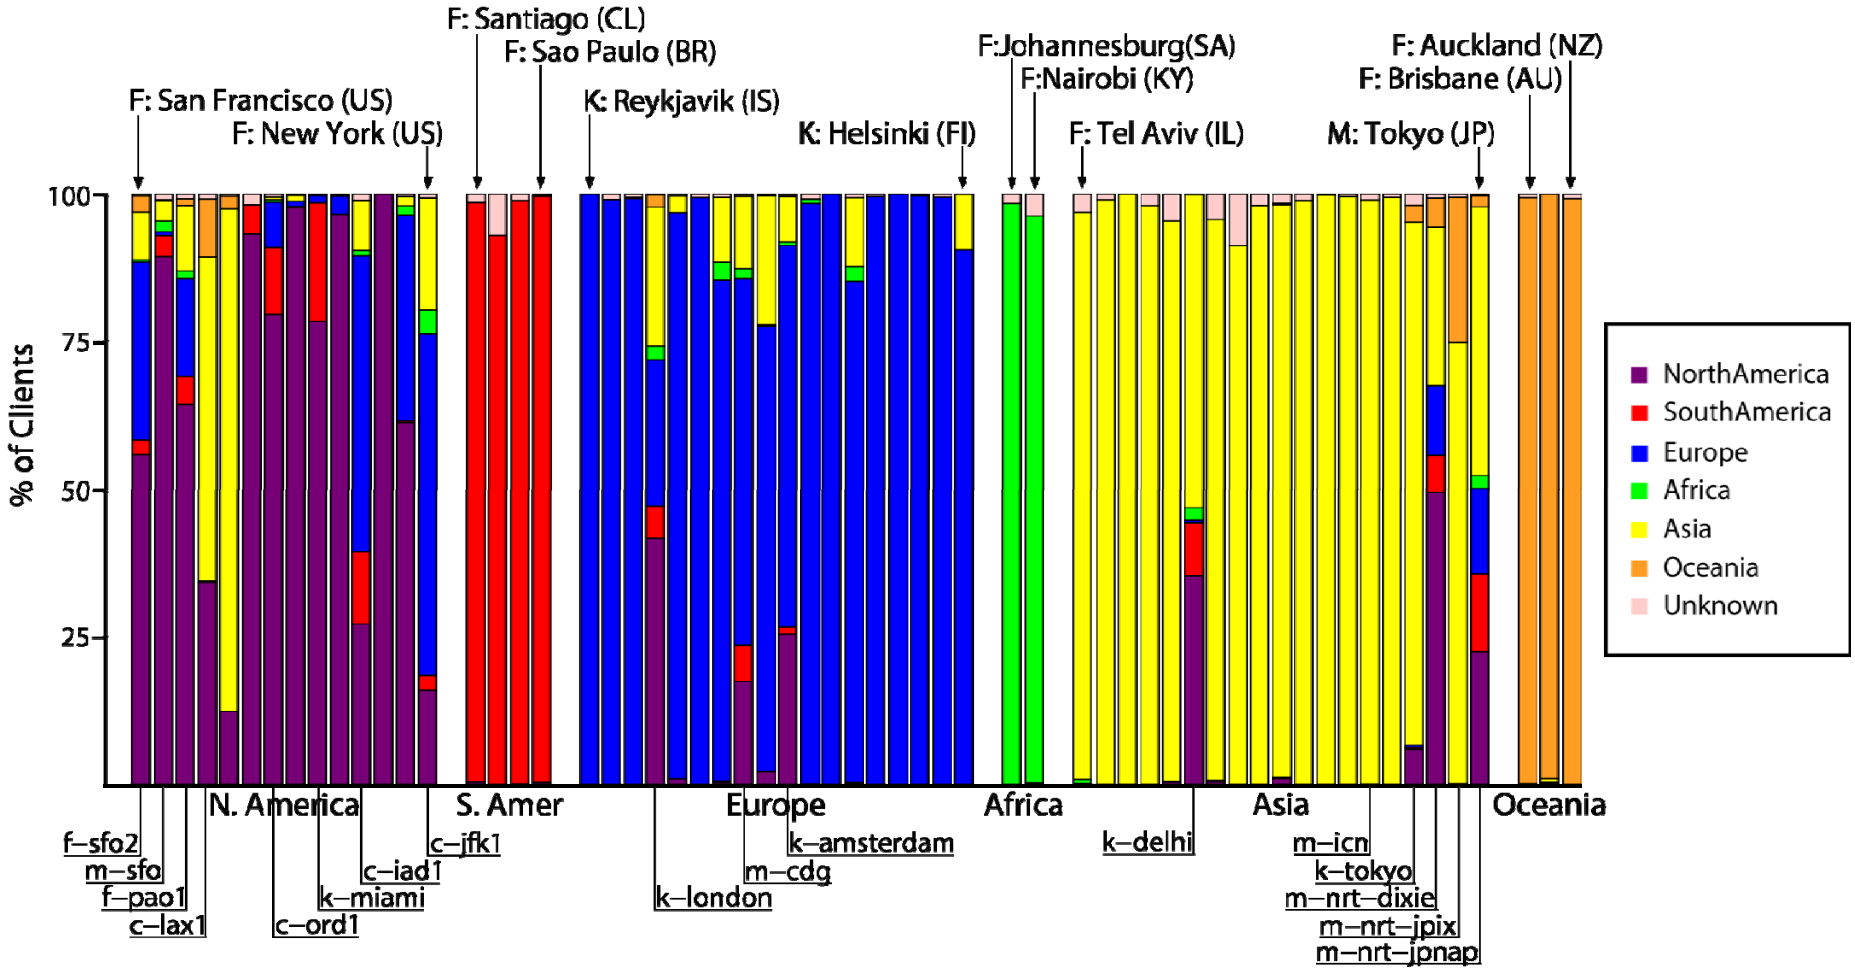

Client Geography

- Local nodes show at least a 91% of clients from the same continent.
  - Exceptions are f-lga (65% from North America, 30% from Europe) and f-lax1 (82.5% from Asia, 14.1% from North America).
- Global nodes
  - F-sfo2 and F-pao1 reduced their Asian clients compared to 2006
    - Perhaps the presence of the local node in Beijing?
    - For Asian clients, F-sfo2 changed for 17% in 2006 to 8% in 2007. In terms of countries, 31.7% from Japan and 19.7 from Rusia to 0.79% and 49.32%.
    - In F-pao1, Asia clients varied from 12.8% to 11%. For China clients, from 36% in 2006 to 2.1% in 2007.

Client Geography

- K-root
  - K-miami and K-tokyo present good correlation with their location.
    - K-miami, 78.4 from North America, 20% from South America.
    - K-tokyo, 88.4 from Asia, 6% from North America.
  - K-london and K-amsterdam have more diverse origins
  - K-delhi has 1/3 of the clients coming from North America.
- M-root
  - All instances receive some amount of traffic from Asia

Query Load

- From a range of 2.54 million clients
  - 438K (17%) sending only one query
    - 61.6% A-query
    - 10% PTR-query
    - 7.1% SOA-query
    - 6% NS-query
    - 4% MX-query
  - 10 sending more than ten million queries.
    - 7 unknown
    - 1 Microsoft Windows NT4
    - 1 Microsoft Windows 2003
    - 1 BIND8

Query Load

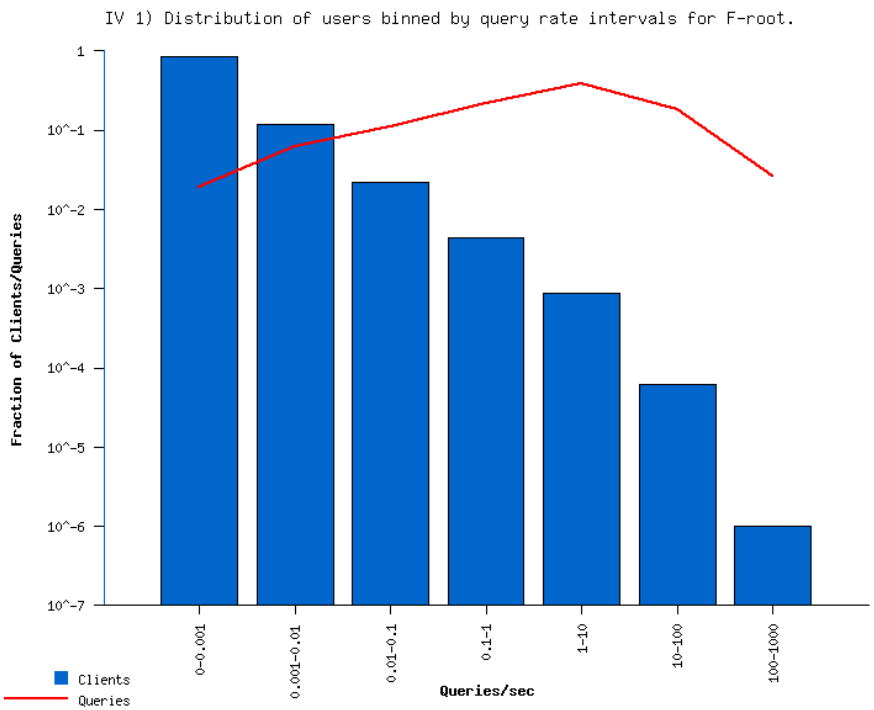
IV 1) Distribution of users binned by query rate intervals for F-root.
Fraction of Clients/Queries
1
10^-1
10^-2
10^-3
10^-4
10^-5
10^-6
10^-7
0-0.001
0.001-0.01
0.01-0.1
0.1-1
1-10
10-100
100-1000
Queries/sec
Clients
Queries

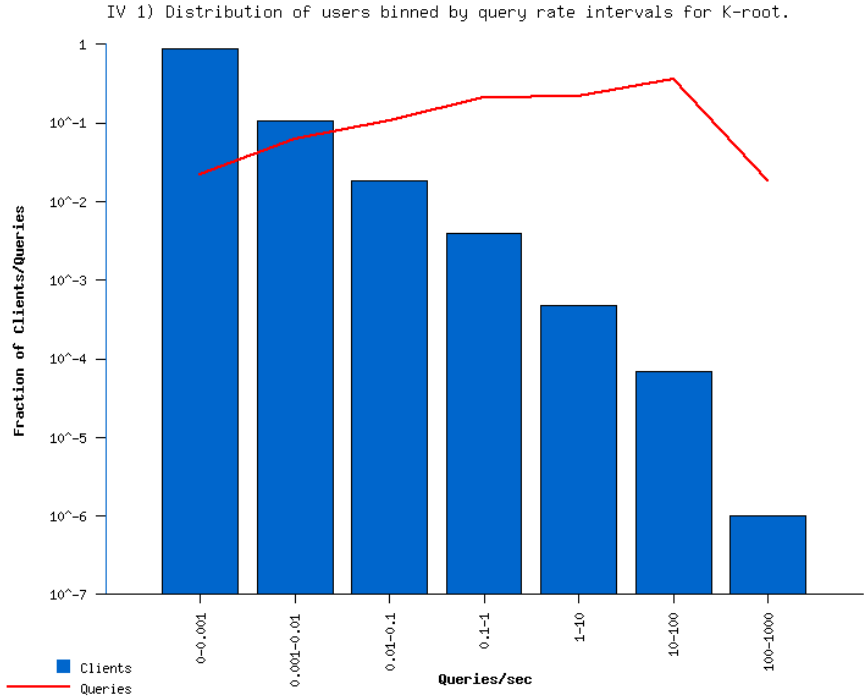
IV 1) Distribution of users binned by query rate intervals for K-root.
Fraction of Clients/Queries
1
10^-1
10^-2
10^-3
10^-4
10^-5
10^-6
10^-7
0-0.001
0.001-0.01
0.01-0.1
0.1-1
1-10
10-100
100-1000
Queries/sec
Clients
Queries

Query Load

- The clients with a query rate less than 1/100 q/s represent 96% of the total client population
  - But only 6% of the load
- For C, F and M, the interval of ‘1-10 q/s’ represents 29%, 39% and 30%.
- For K root, the interval of ’10-100 q/s’ represents 38.7% of the queries

Query Load

- The same categories per query rate
- Each column show the fraction of queries per type
- The two rightmost columns (higher query rates) present different behavior
  - M-root
    - 25% A queries
    - 50% PTR queries
    - 13% MX queries
  - K-root
    - 55% A queries
    - 35% SOA queries
    - 10% PTR queries

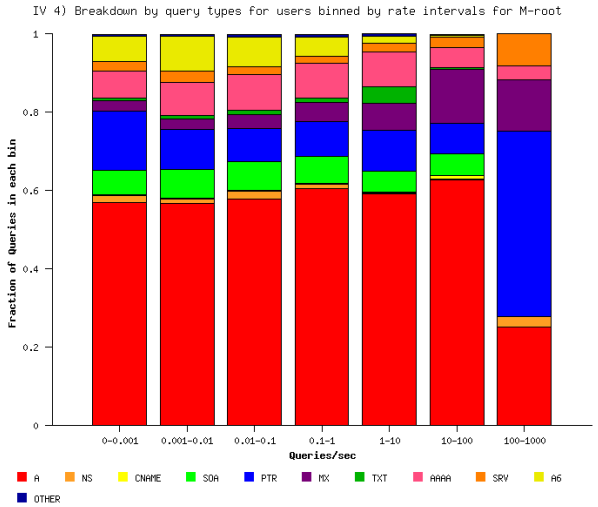

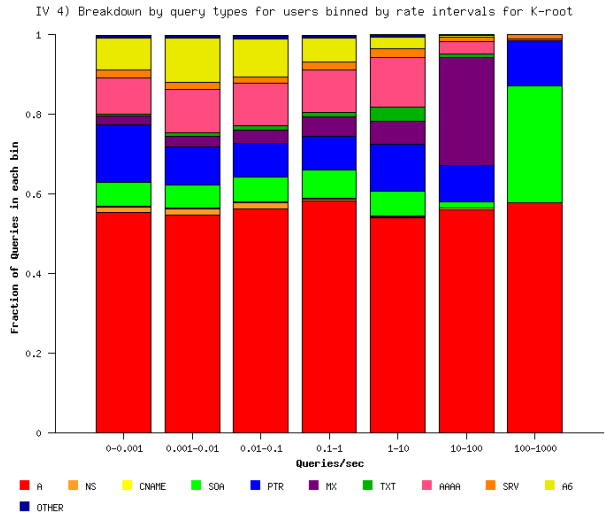

Query Load

- C-root and F-root similar
  - 90% A queries
  - 5% PTR queries

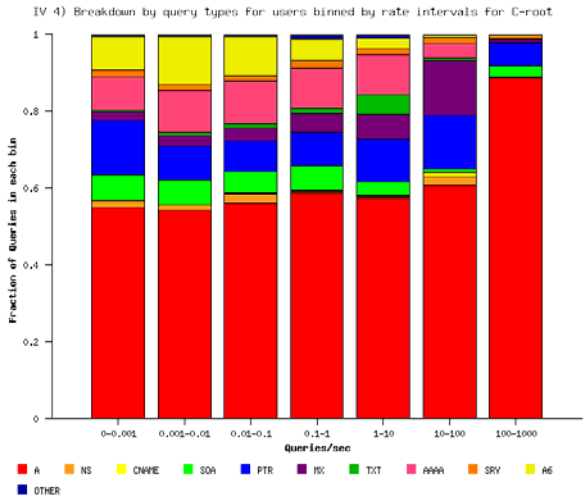

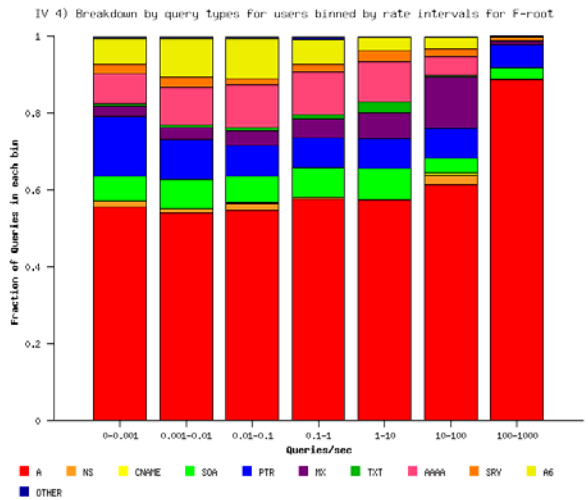

Client affinity

- 39% queried one cloud
- 28.3% queried all four
- Instance switch
  - C-root: 0.9%
  - F-root: 0.7%
  - K-root: 1.2%
  - M-root: 1.9%

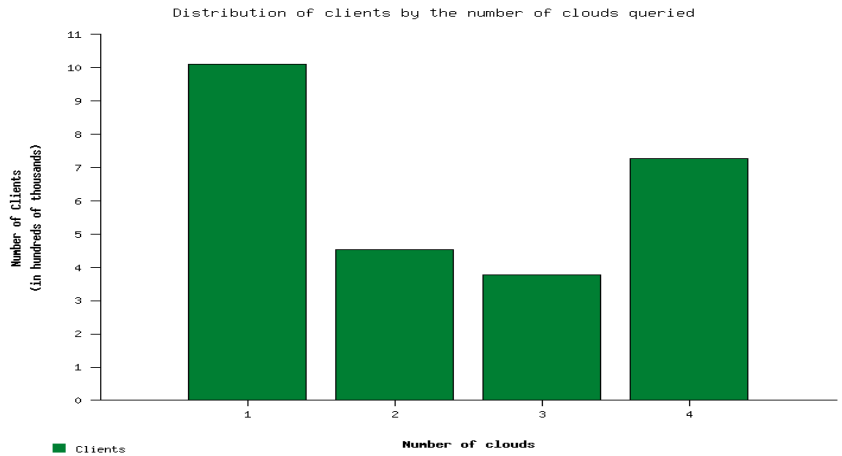

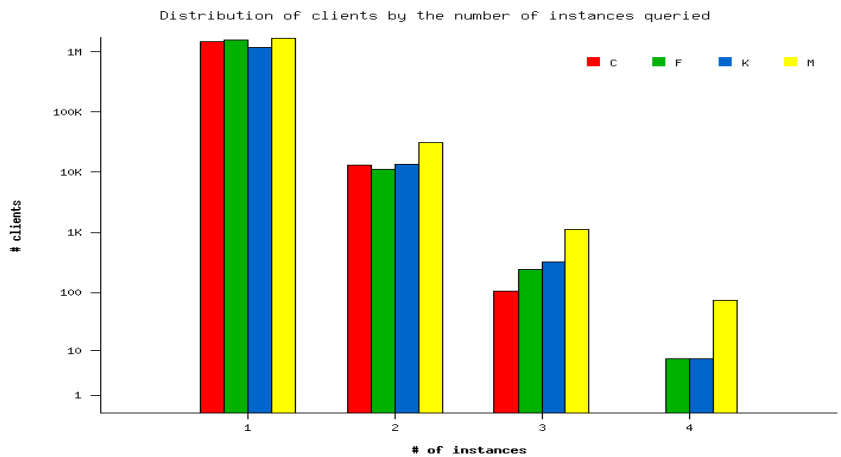

Invalid queries

- Methodology
  - Updating the results from a paper of 2003.
  - Nine categories of invalid traffic
  - Evaluation one by one
  - For the last three, the analysis was per source address
    - Requires filter/split traffic per source
    - Using a sample (7.5% clients for each cloud)
      - Currently 213142 unique clients
      - The goal is reach 10% of clients per cloud

Invalid queries

- Unused query class: Any class not in IN, CHAOS, HESIOD, NONE or ANY
- A-for-A: A-type query for a name is already a IPv4 Address
  - <IN, A, 192.16.3.0>
- Invalid TLD: a query for a name with an invalid TLD
- Non-printable characters: a query for a name with characters not in [A-Z0-9\-] list
- Queries with '_': Special category for the invalid but widely used character.
- RFC 1918 PTR: a PTR query for an IPv4 address in the private space
- Identical queries: a query with the same class, type, name and id (during the 24 hours period)
- Repeated queries: a query with the same class, type and name
- Referral-not-cached: a query seen with a referral previously given.
  - If a client sent <IN, A, www.example.net> and later <IN, NS, ripe.net> the second query counts as "referral-not-cached" because a referral to "net" nameservers was answered.
  - A tolerance parameter of 2 seconds was included on this analysis
  - Root servers are authoritative for .arpa, .in-addr.arpa and root-servers.net zones, were included as special cases.

Invalid queries

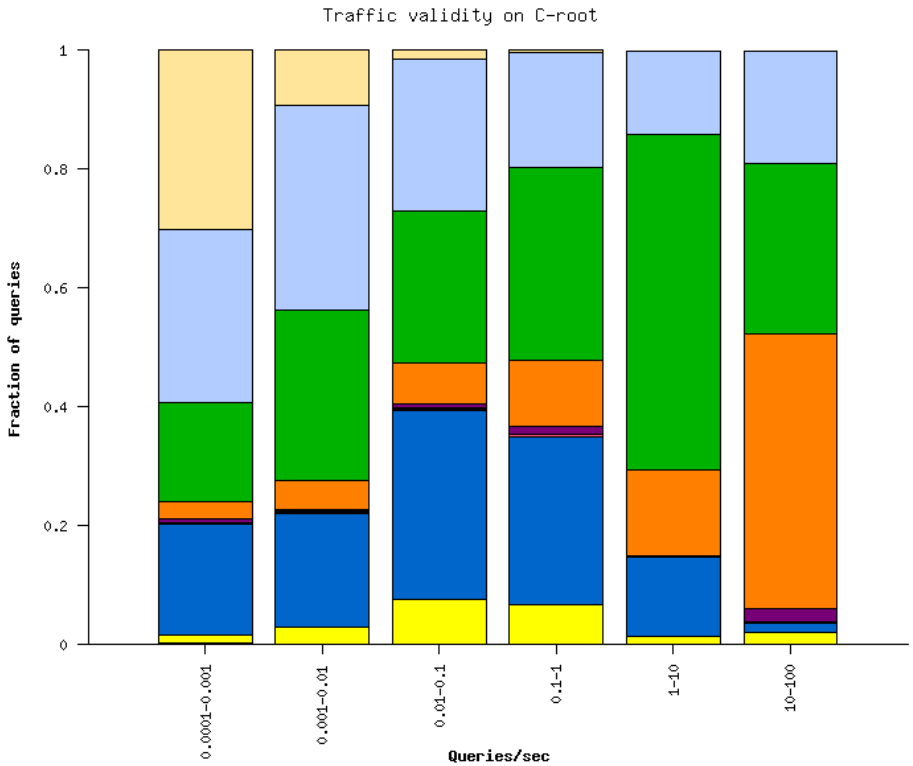

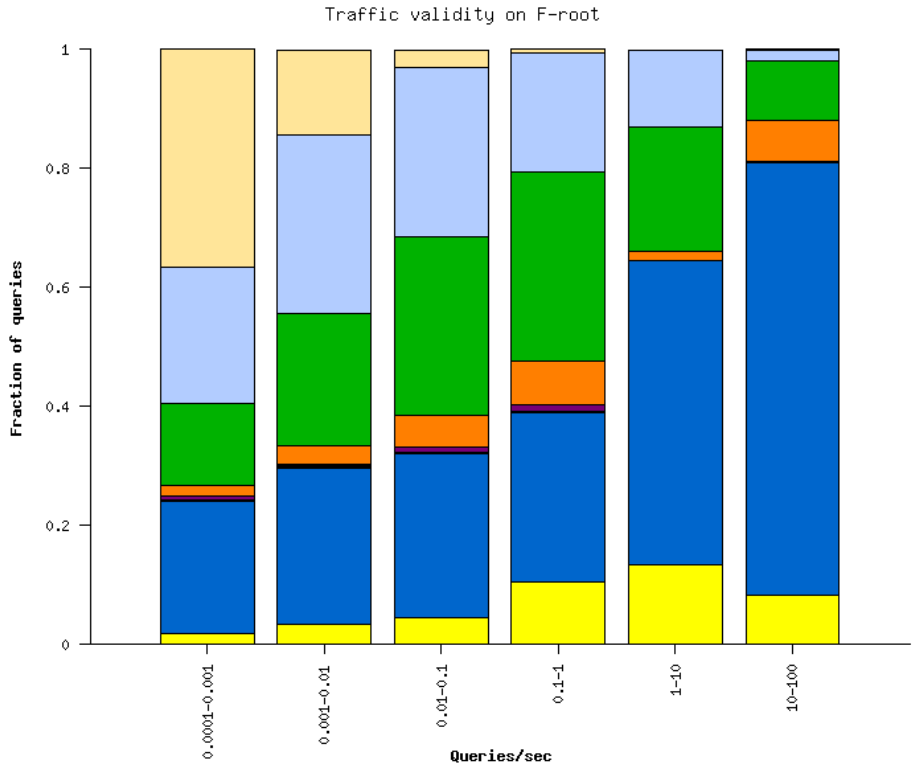

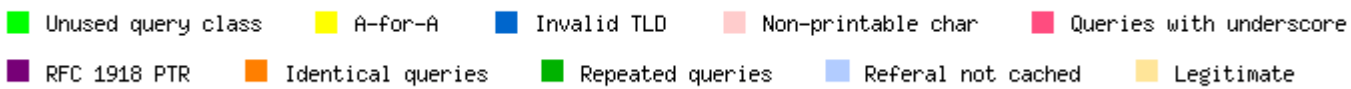

Invalid queries

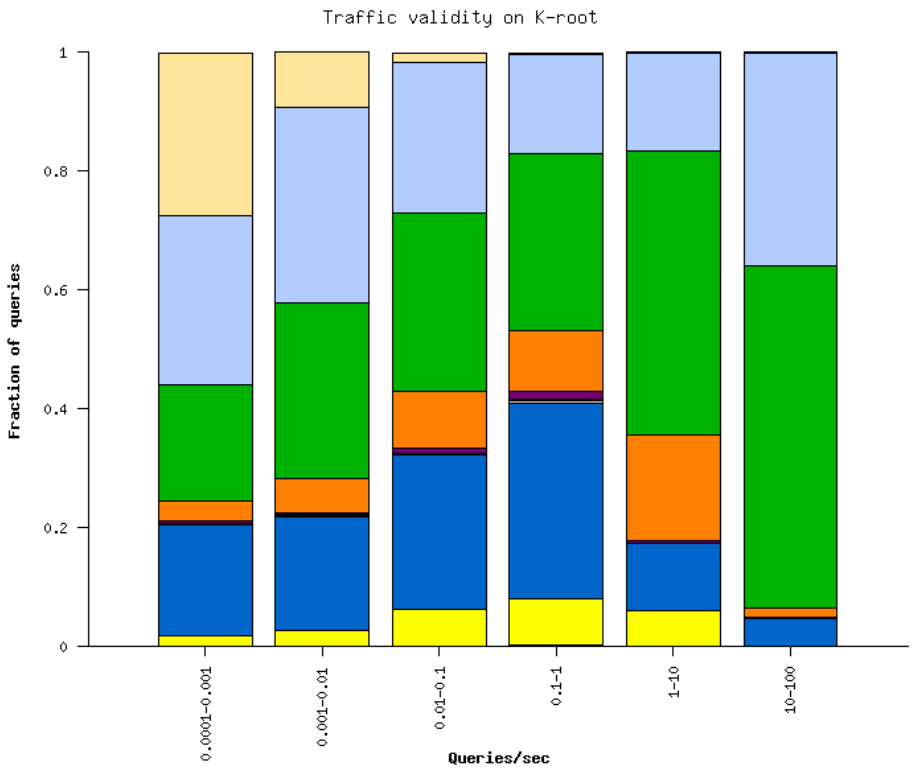

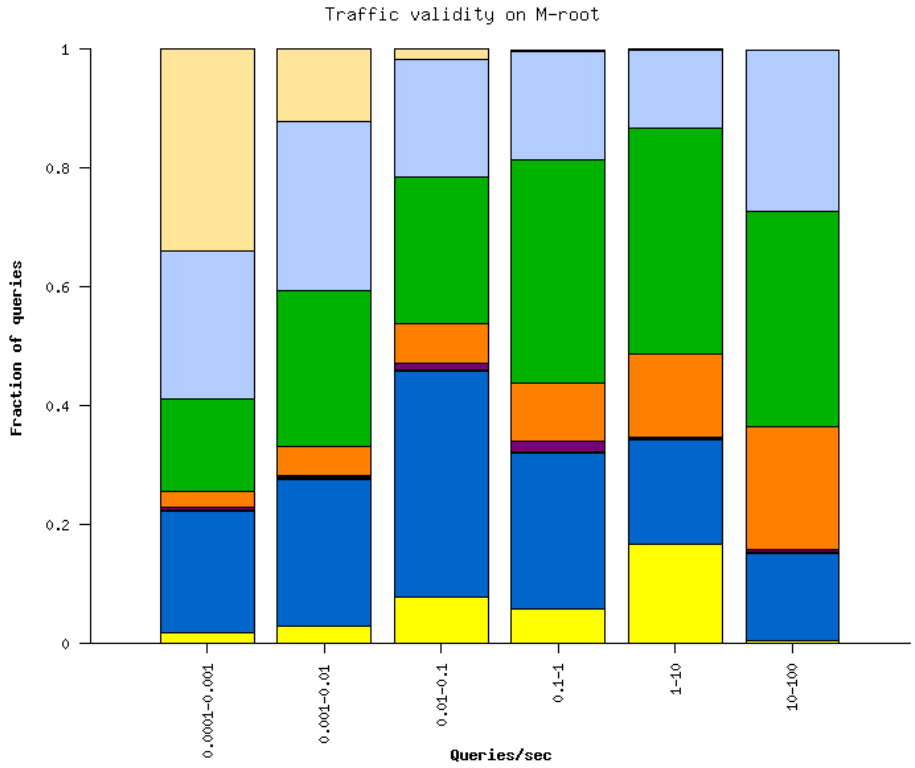

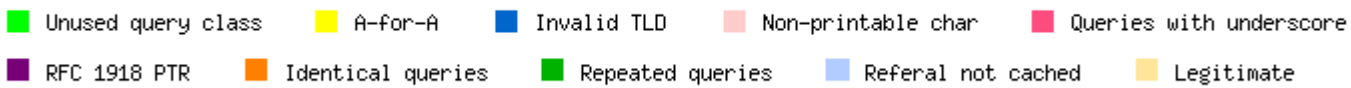

Invalid queries

Category	C-root	F-root	K-root	M-root	Total Sample	Total
Unknown Class	0.09	0.01	0.02	0.05	0.03	0.08
A-for-A	4.65	11.57	2.18	8.74	6.64	7.02
Invalid TLD	19.15	46.79	10.01	20.96	24.72	24.73
Non-printable Character	0.05	0.03	0.13	0.10	0.08	0.53
Queries with '_'	0.13	0.07	0.15	0.13	0.12	0.23
RFC-1918 PTR	0.84	0.28	0.22	0.74	0.44	0.67
Identical Queries	15.41	3.73	4.78	12.51	7.71	N/A
Repeated Queries	37.99	20.11	50.20	32.61	35.73	N/A
Referral not Cached	18.69	15.25	30.55	21.22	22.07	N/A
Legitimate Queries	3.01	2.17	1.76	2.95	2.46	N/A

Invalid queries

- Common invalid TLD's

TLD	Percentage of queries
local	20.29
localhost	8.92
domain	3.15
invalid	2.43
lan	2.06
belkin	1.76
home	1.30
localdomain	1.29
wpad	0.74
txt	0.74

Invalid queries

- Refinements
  - We explored how many repeated queries could be originated by BIND9 ‘glue record refresh’.
  - We found at most 1% of the total queries from sample could be associated to that process
  - Still adjusting analysis to find common TTL parameters (and differentiate from dual stack clients sending A/AAAA queries for the same name).

Invalid queries

- Comparing with data from ORSN (Open Root Server Network)

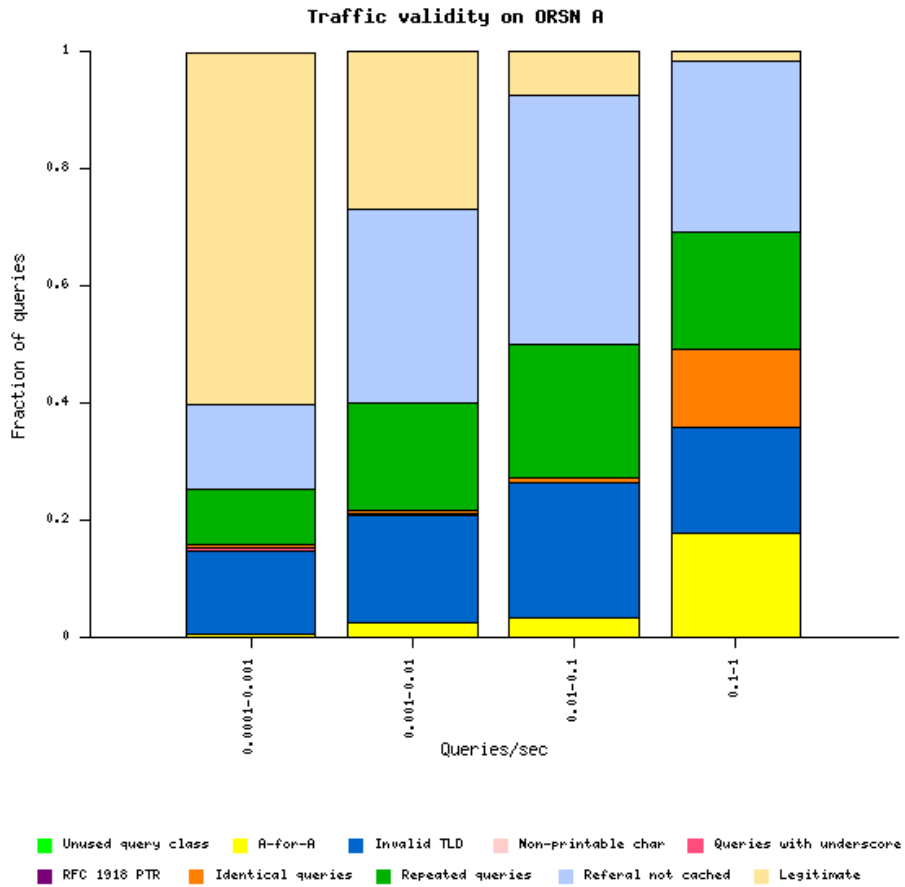

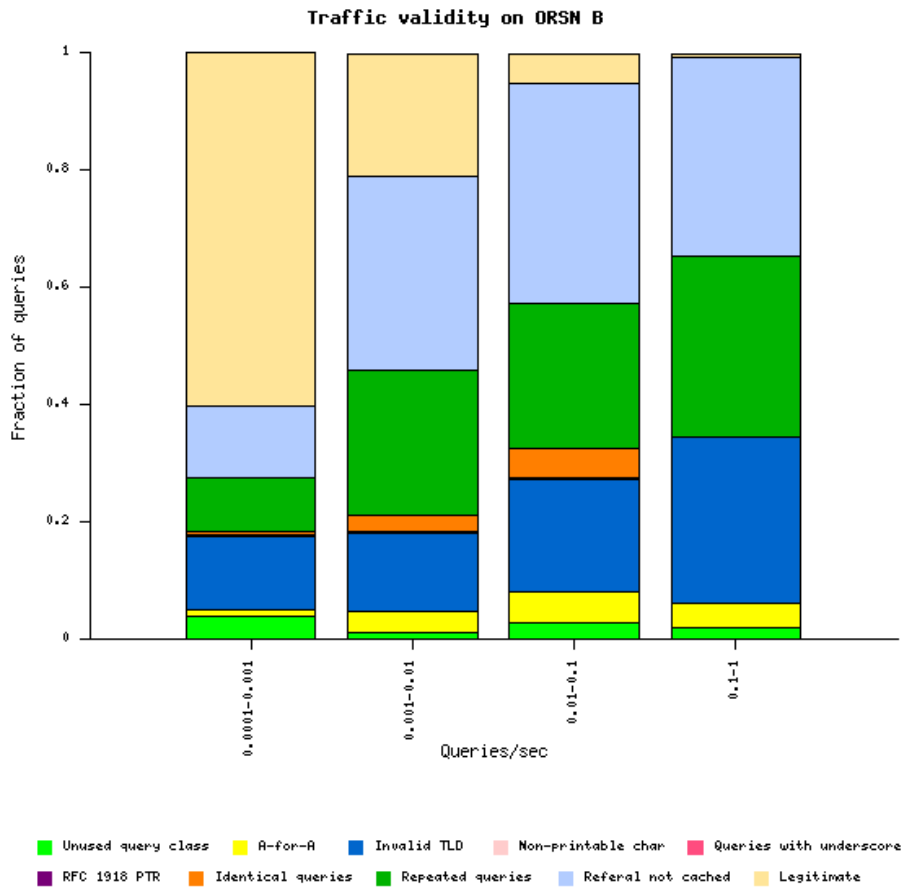

Source port

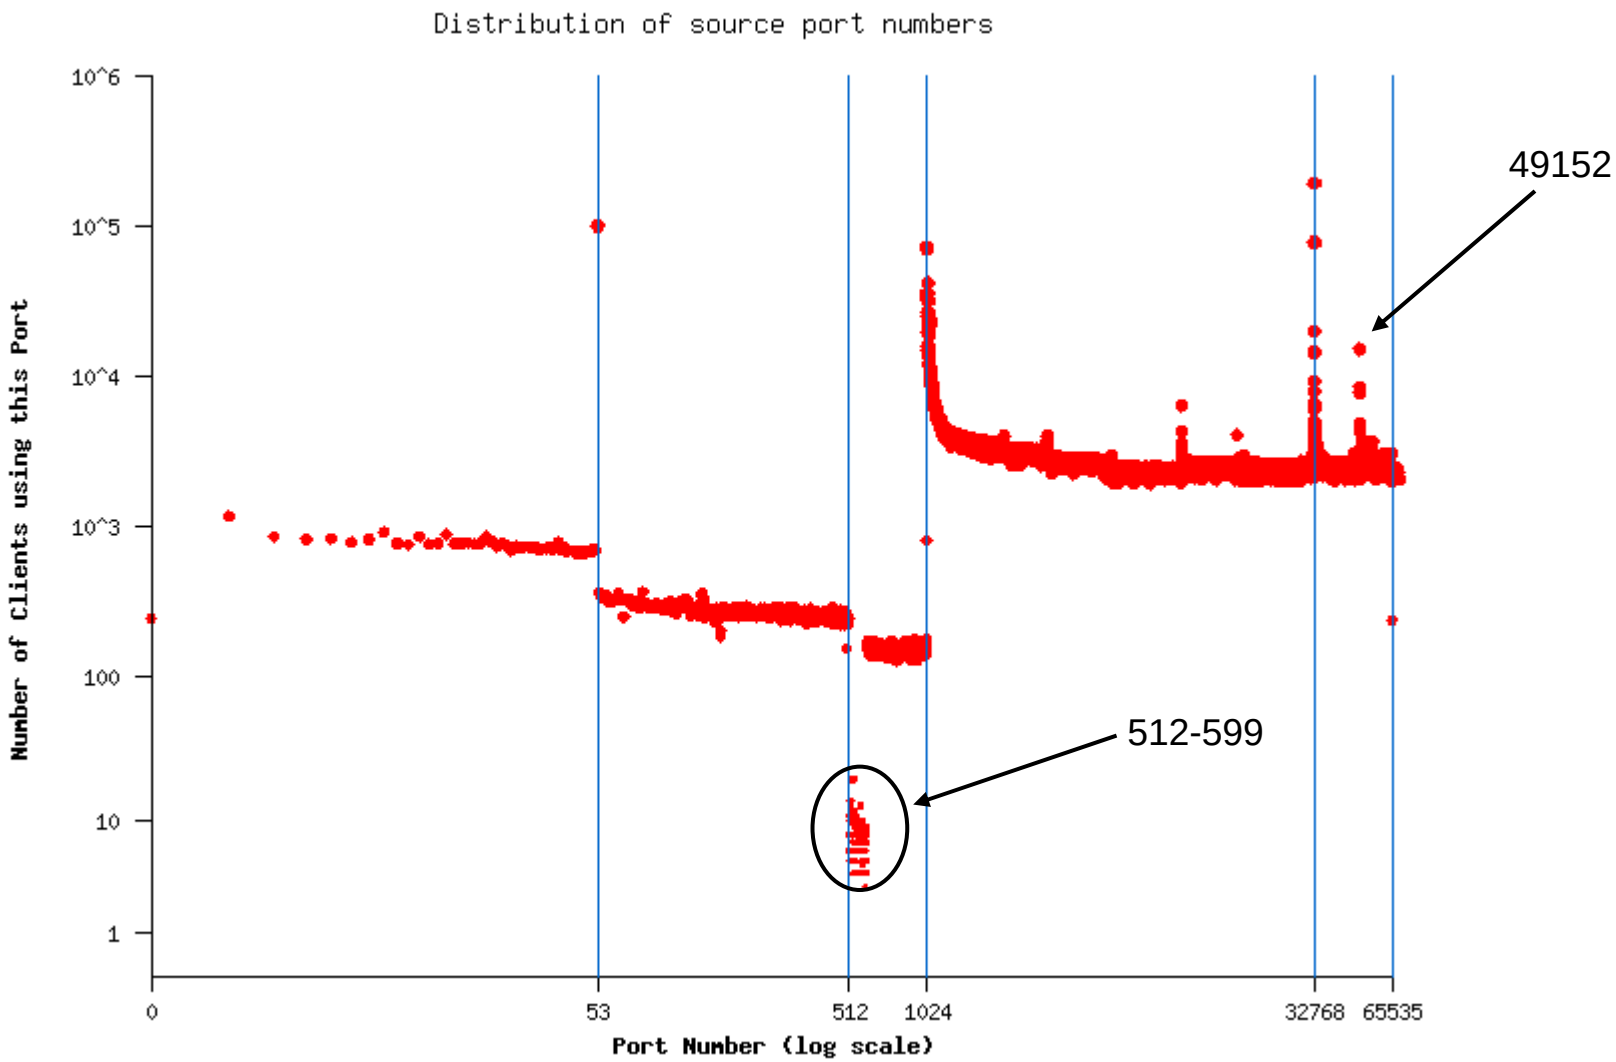
Distribution of source port numbers
Number of Clients using this Port
10^6
10^5
10^4
10^3
100
10
1
0
53
512
1024
32768
65535
Port Number (log scale)
49152
512-599

Source port

- The idea was suggested on last OARC meeting
- Presents the distribution of client by the source ports on queries received on root servers
- Use of port 0, 53, 512, 1024, 32768.
- 49152 = (32768 + 65536) / 2 (?!)
- Use of “privileged port range”
  - BSD kernel has settings to use ports starting at 600
- What about trying doing some O.S. fingerprinting using IP parameters and port range?

EDNS

- For every query with a peudo-RR OPT in Additional section, calculated EDNS version and EDNS buffer size.
- Found some clients sending more than one buffer size (not included on the graphs)
- The EDNS version graphs are not new
  - Already supported by DSC
- The buffer size is a little more interesting.
  - Shows, for example, queries with buffer of 512 bytes (but not present on per client graph)
  - Explained by Mark Andrews: BIND9 falling back to 512 bytes.

EDNS

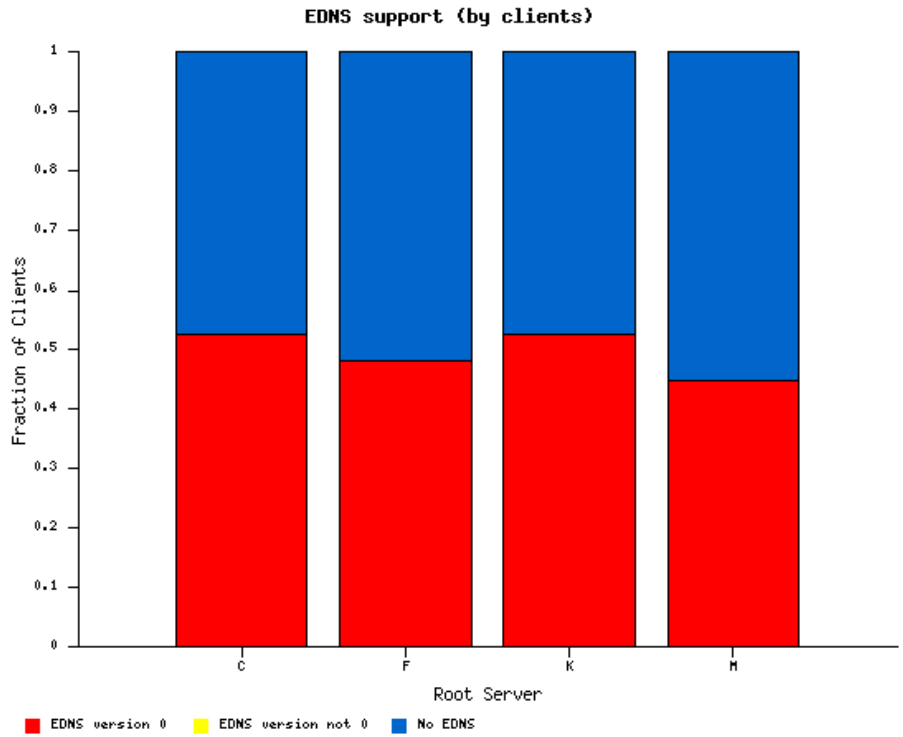
EDNS support (by clients)
Fraction of Clients
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
C
F
K
M
Root Server
EDNS version 0
EDNS version not 0
No EDNS

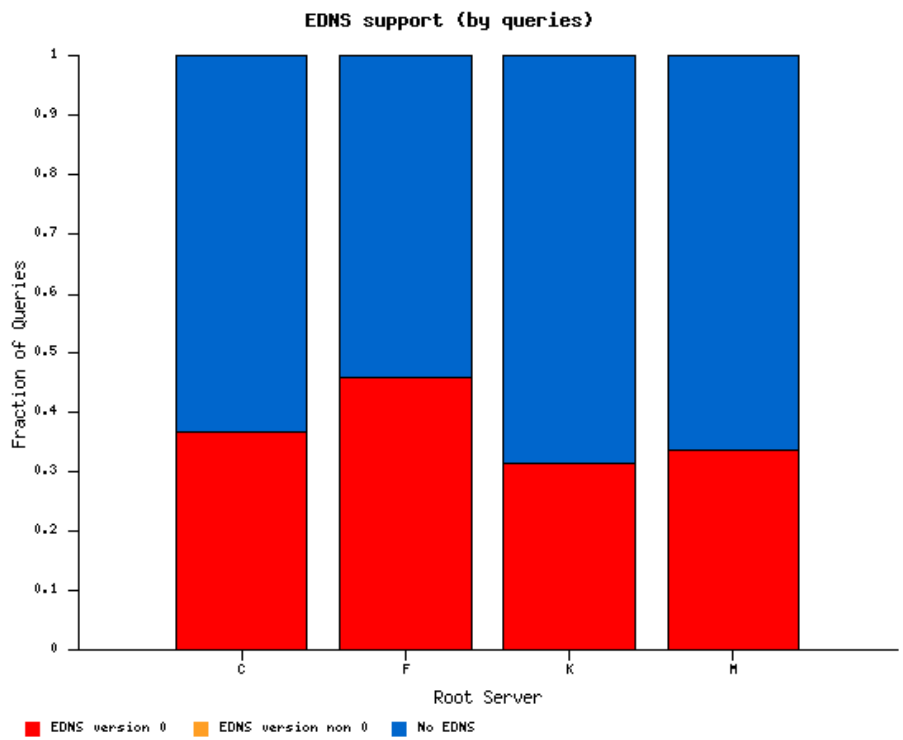
EDNS support (by queries)
Fraction of Queries
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
C
F
K
M
Root Server
EDNS version 0
EDNS version non 0
No EDNS

EDNS

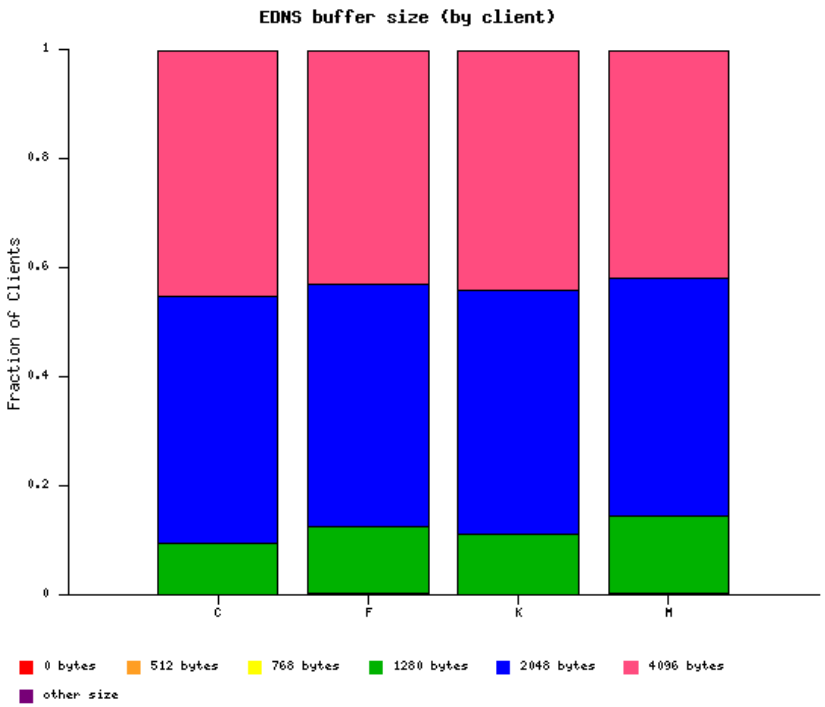
EDNS buffer size (by client)
Fraction of Clients
1
0.8
0.6
0.4
0.2
0
C
F
K
M
0 bytes
512 bytes
768 bytes
1280 bytes
2048 bytes
4096 bytes
other size

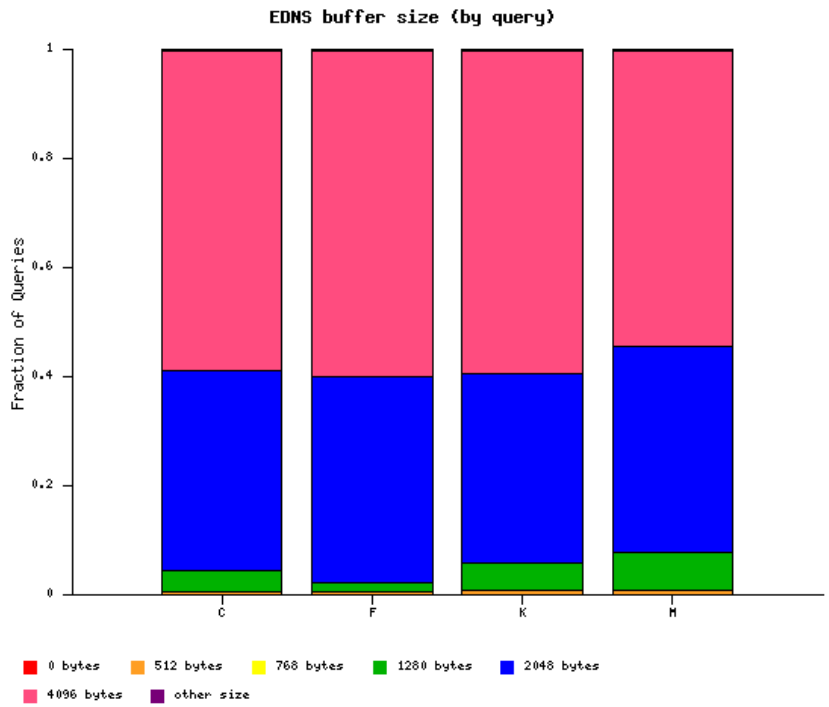
EDNS buffer size (by query)
Fraction of Queries
1
0.8
0.6
0.4
0.2
0
C
F
K
M
0 bytes
512 bytes
768 bytes
1280 bytes
2048 bytes
4096 bytes
other size

Conclusions

- Reduced number of node switching compared with 2006
  - Previously seen on J-root analysis
- Still low TCP traffic
- After 4 years, the root still sees the same amount of trash
  - Should be reasonable/effective to take measures about this, from education to punishment?