
2007 Day In The Life DNS Root Server Analysis


Duane Wessels
The Measurement Factory/CAIDA
&
Sebastian Castro
NIC Chile/CAIDA

RSSAC
December 2, 2007

DITL 2007

- Day In The Life of The Internet. Okay, two days.
- 48 hour period: Jan 9 00:00:00 to Jan 10 23:59:59 UTC
- Primary focus is DNS and root servers, but other data was collected as well.
- We have data from C-, F-, K-, and M-roots, which is the subject of this presentation. E-root traces were lost.
- Also data from two ORSN servers.
- Data is 740 GB compressed pcap files.
- 10,000,000,000 DNS queries.
- This analysis is based on a 24-hour subset of collected data.

Coverage Chart



Terminology


- Server: a collection of DNS nameservers operating under the same IP address.
 - c.root-servers.net is a server
- Instance: an anycast instance of a server.
 - k-milan is an instance of k.root-servers.net.

Load-balanced nodes are combined into a single instance.

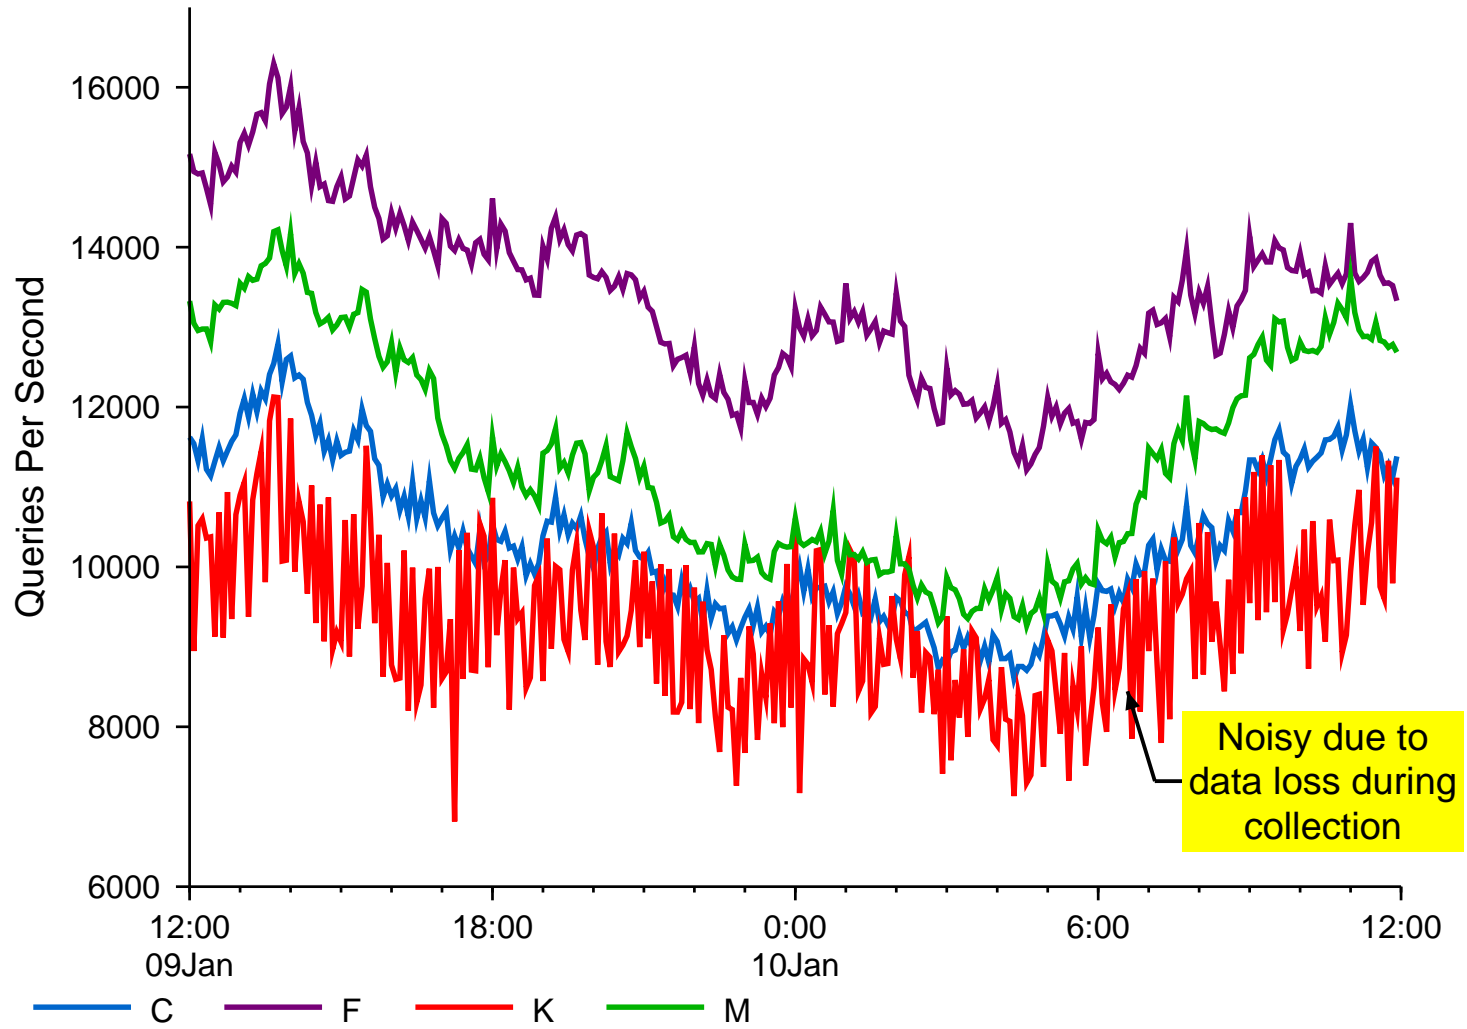
- c-lax1a and c-lax1b are load-balanced members of the c-root LAX instance.
- Client: an IP address sending DNS queries.
 - We don't try to distinguish multiple (NAT-ed) clients sharing a single IP address.
 - Or cases where an address is assigned to multiple hosts in different locations

Basics


II 1) Average rates of requests.




II 1) Average rates of requests.




II 2) The average number of clients per second seen at each instance.




II 2) Number of clients per second seen at each M-root instance.




IV 1) Distribution of users binned by query rate intervals for C-root.



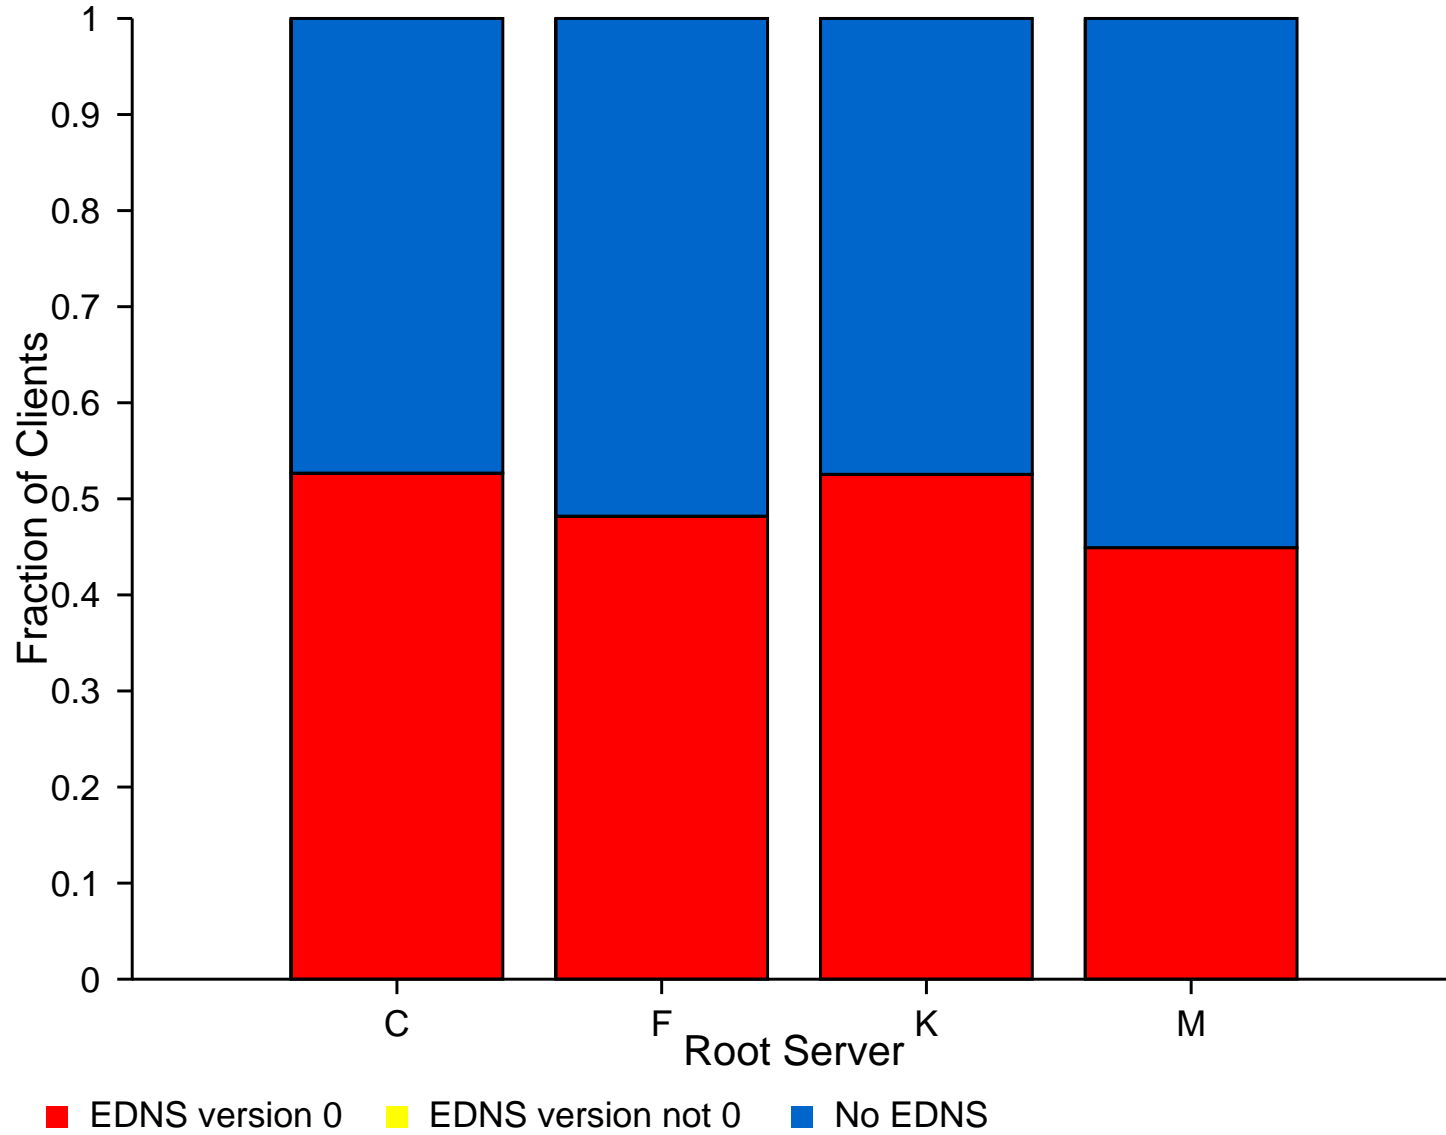
IV 3) Breakdown by query types




IV 4) Breakdown by query types for users binned by rate intervals for M-root




VII 2) EDNS support (by clients)



EDNS buffer size (by client)




Distribution of source port numbers




Routing and Anycast

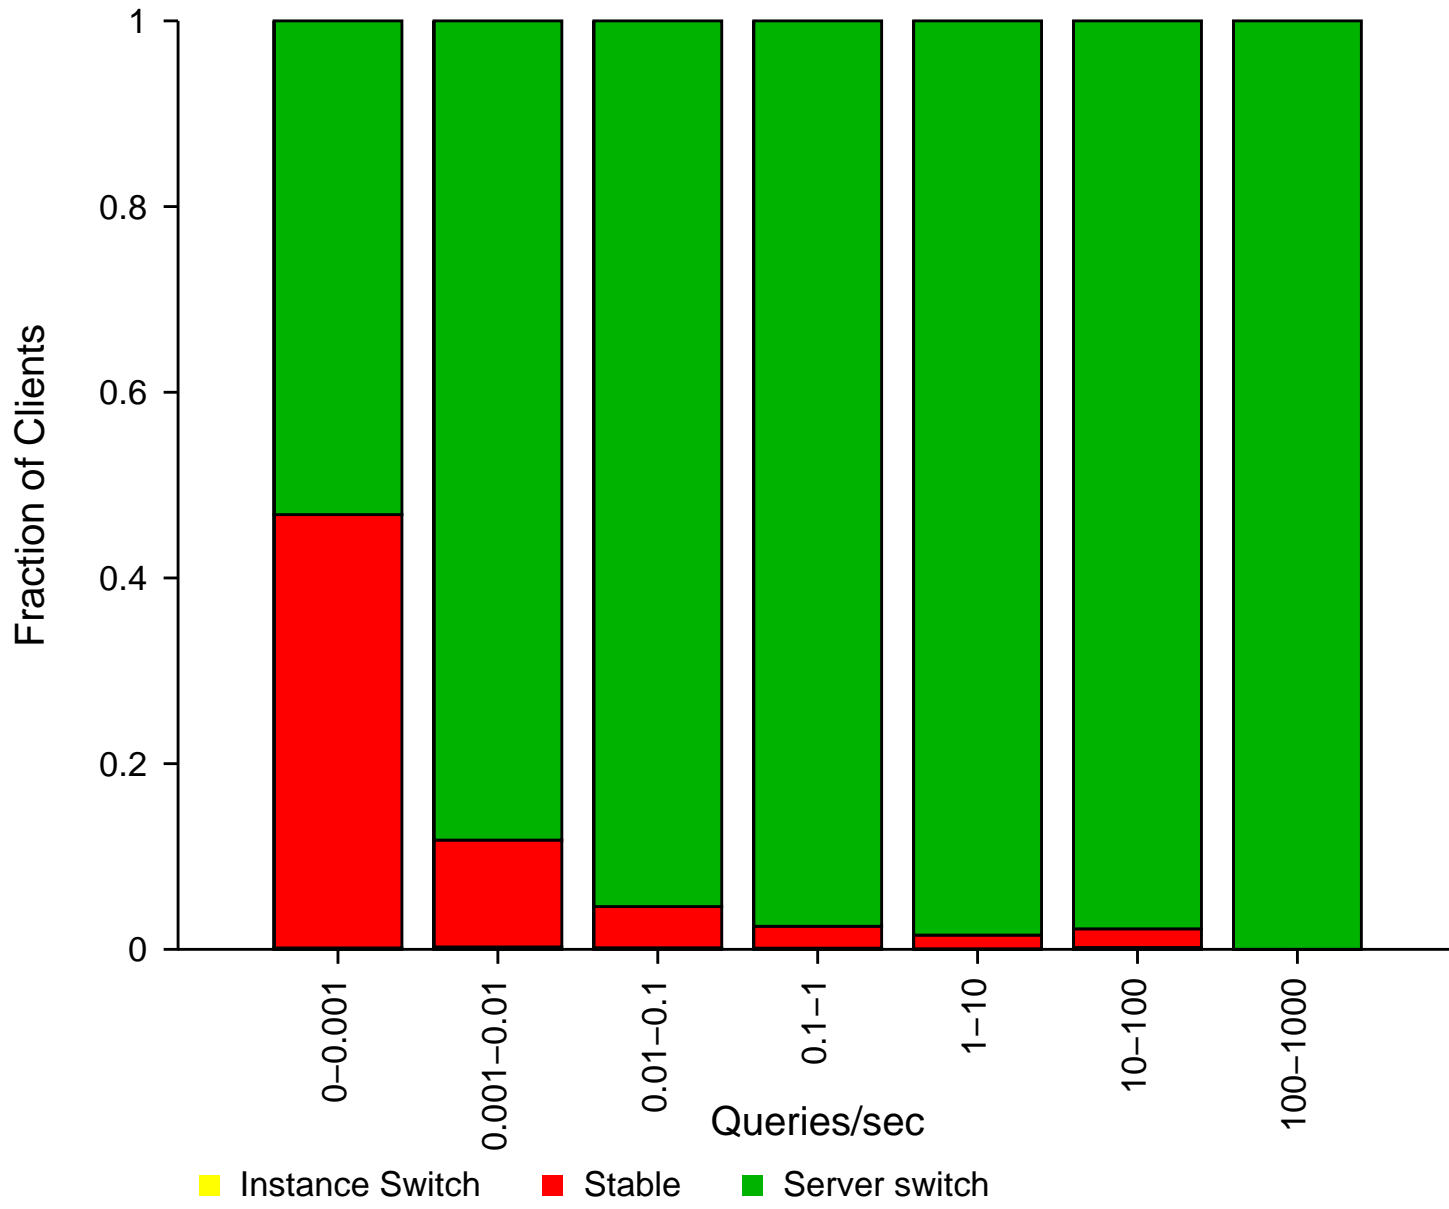
II 3) Topological coverage by ASes.




II 4) Topological coverage by prefixes.




Client classification by stability



Percentage of the clients switching instances

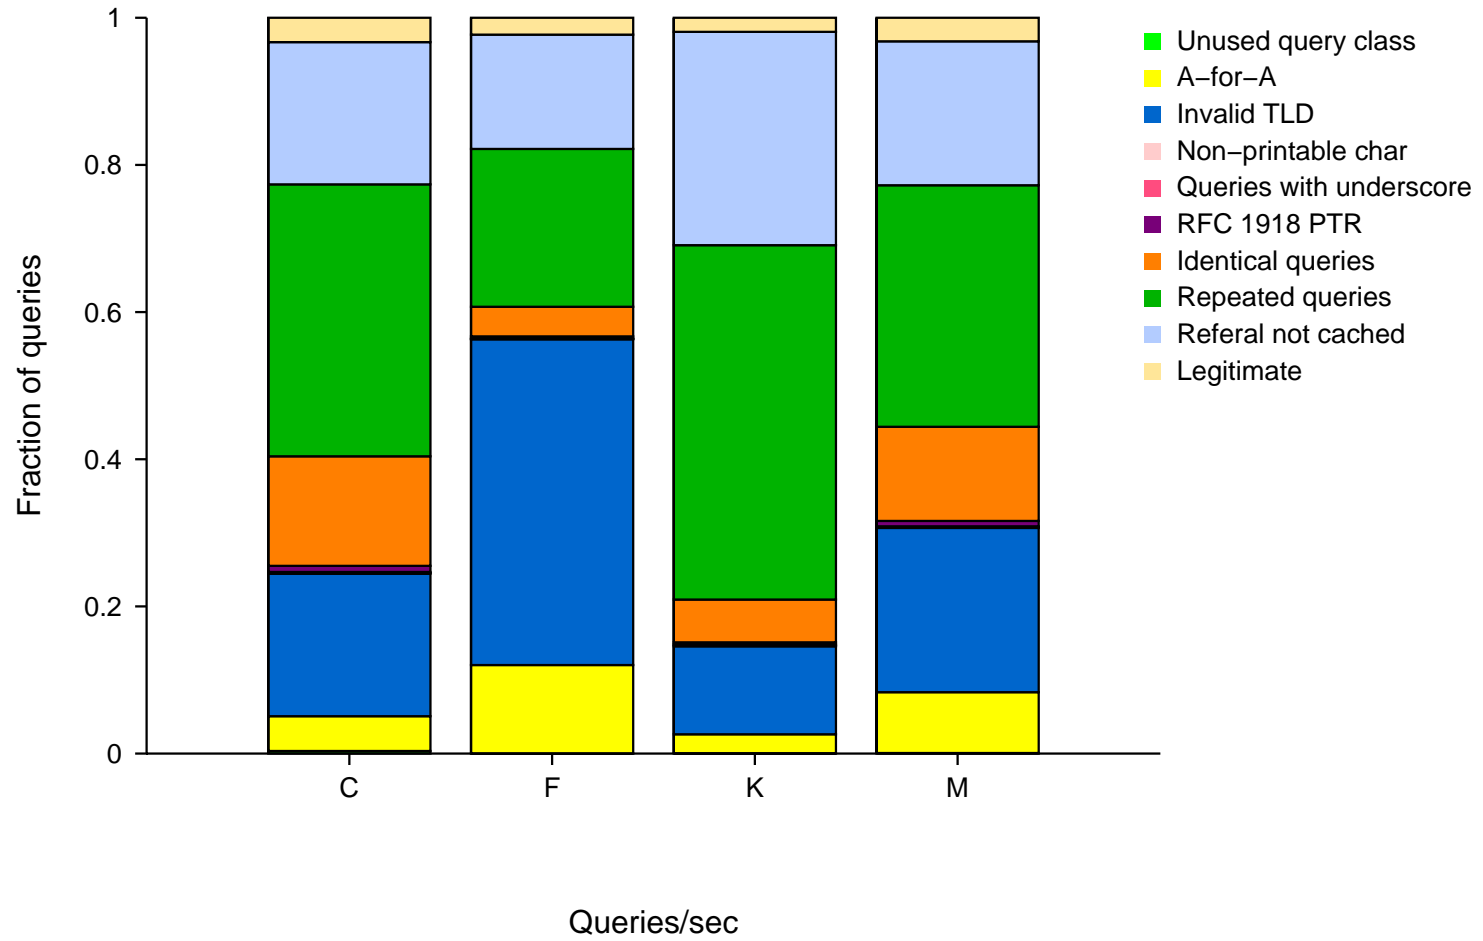


Distribution of clients by the number of instances queried

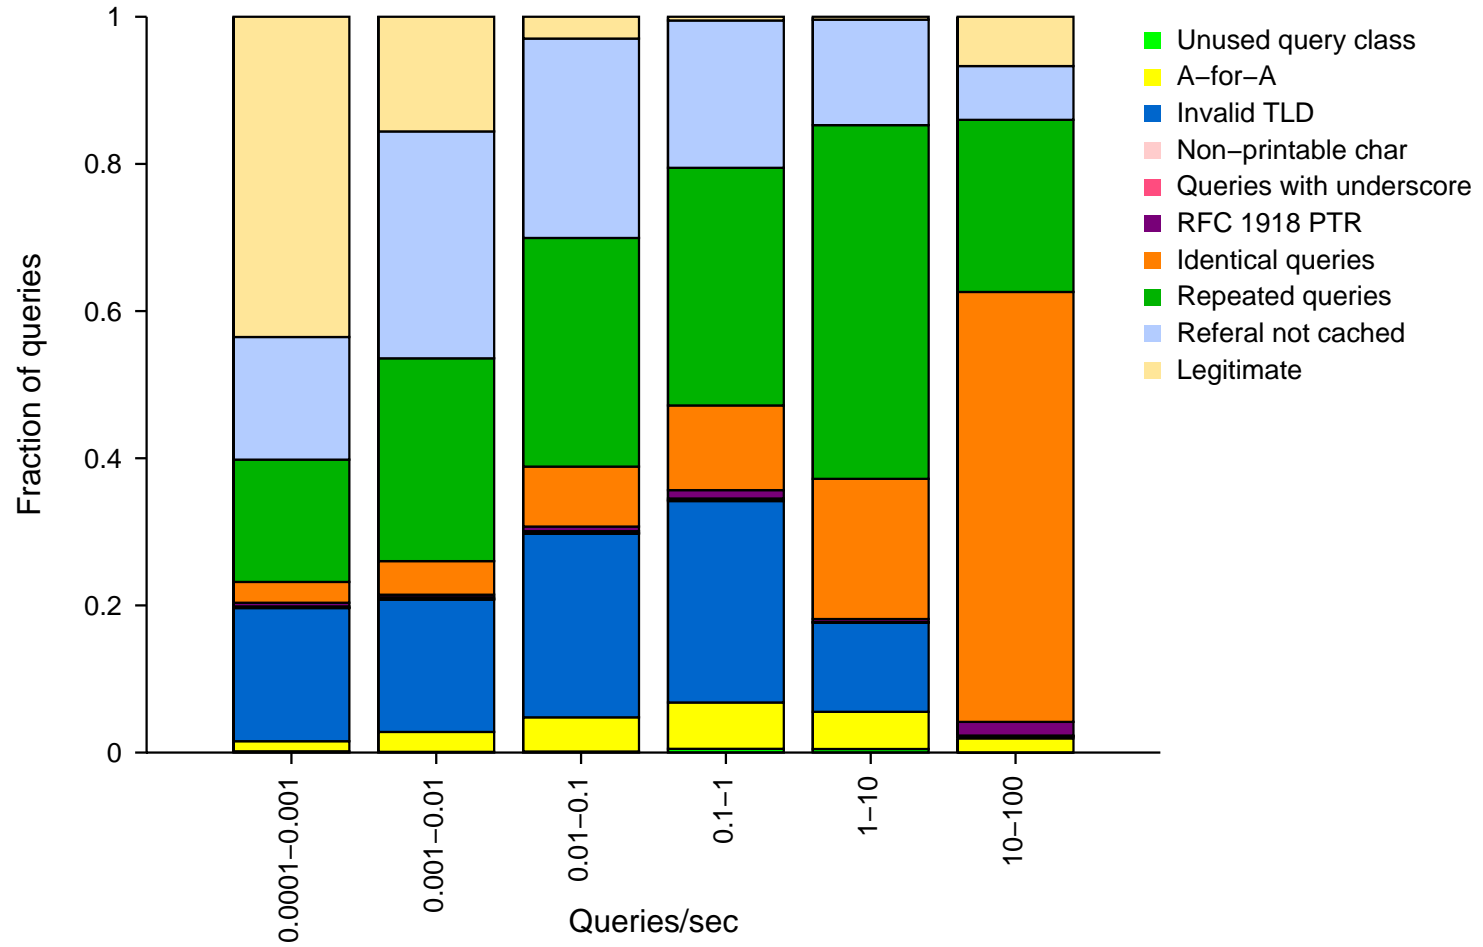


Query Validity

Query Validity



Query Validity on C-root




Comparison to F-root in 2002

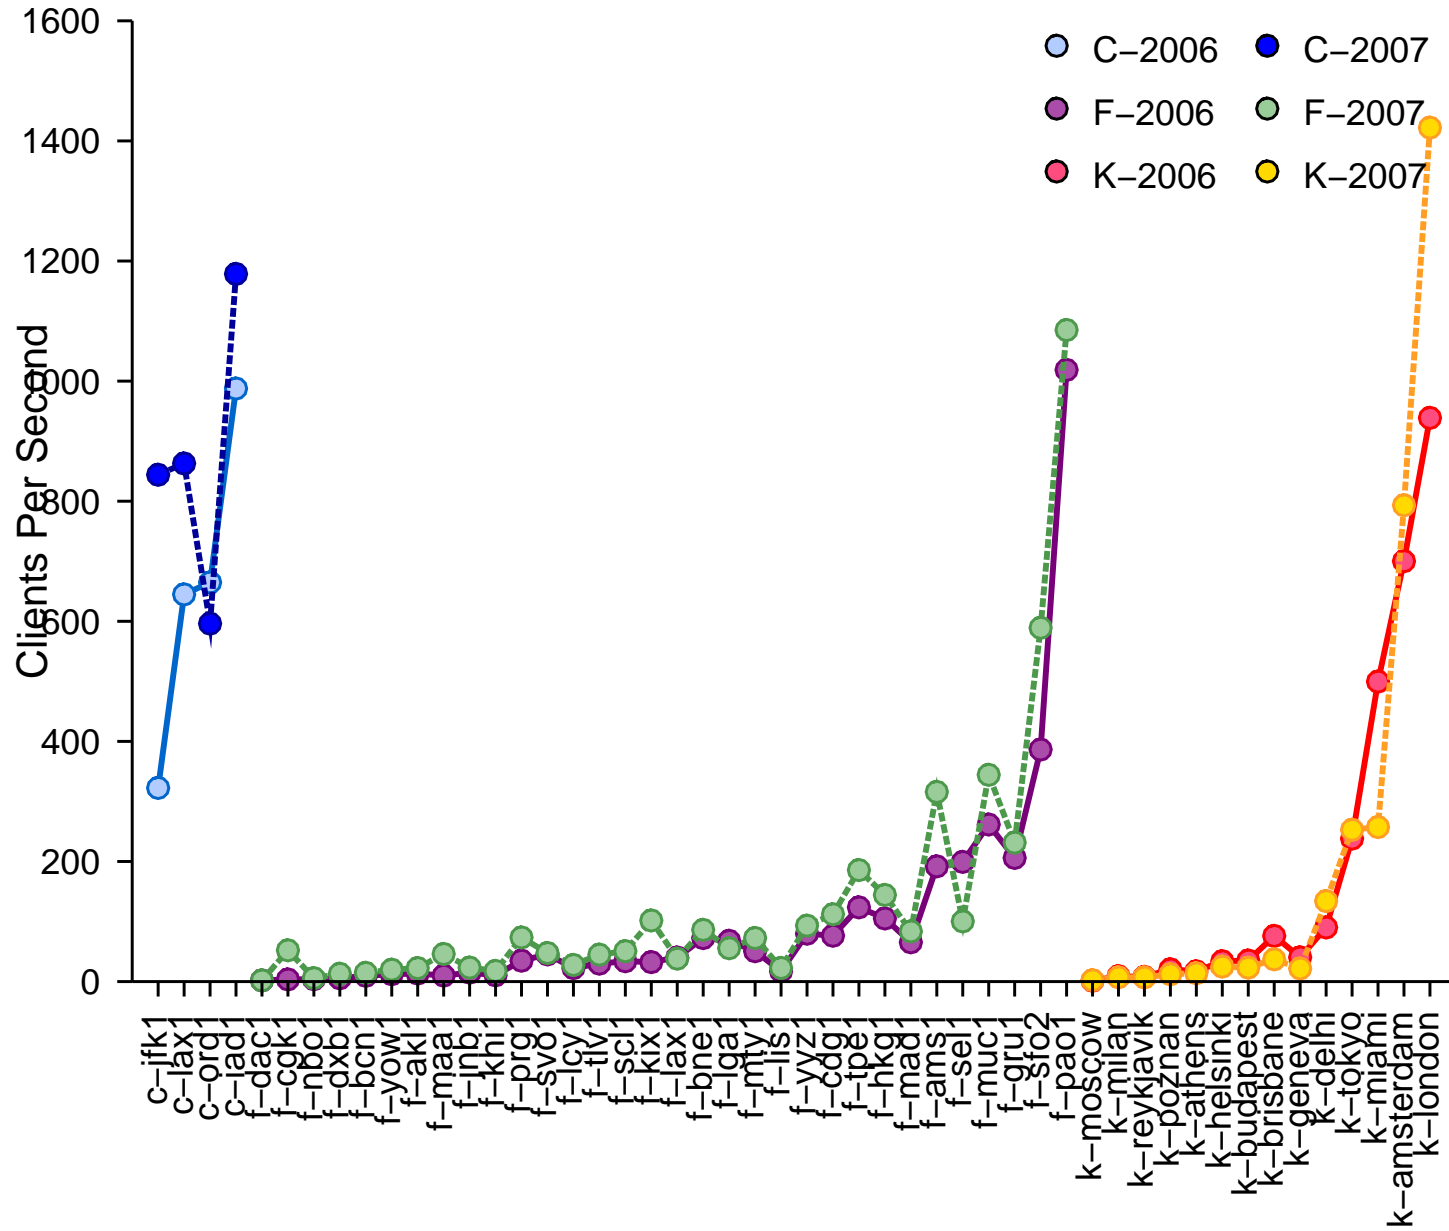
| Category | 2002 % | 2007 % |
|-----------------------|-------------|-------------|
| Unused Query Class | .024 | .018 |
| A for A | 7.03 | 12.0 |
| Unknown TLD | 12.5 | 44.3 |
| Nonprintable in query | 1.94 | .031 |
| RFC1918 PTR | 1.61 | .284 |
| Identical Query | 25.4 | 4.03 |
| Repeated Query | 44.9 | 21.5 |
| Referral Not Cached | 4.36 | 15.5 |
| Legitimate | 2.15 | 2.30 |

Comparison to 2006


Mean query rate on 2006 and 2007




Mean client rate on 2006 and 2007



Breakdown by query types



Distribution of source port numbers



More Info

- Raw DITL traces available through DNS-OARC
- <http://www.caida.org/projects/ditl/summary-2007-01/>
 - Describes additional, non-DNS-rootserver data collected for DITL'07.
- DITL'08 collection to take place late February or early March.
- DITL'08 Planning at CAIDA/WIDE Workshop January 2008, Hawaii

The End