An analysis of DITL root data and comparison with JP data

Kazunori Fujiwara, JPRS

<fujiwara@jprs.co.jp>

DNS-OARC 2013 Fall Workshop

Abstract

- The number and characteristics of full resolvers are presumed in analyzing DITL data and JP packet capture data.
- This report presents number of IP addresses which send:
  - root DNSKEY queries
  - EDNS0 queries, DO=1 queries
  - non-exististent name queries
  - updates, JP domainname queries and others.
- Then, it compares root data and JP data.

Datasets and analysis method

DNS-OARC Root Datasets

- "A Day in the Life of the Internet" (DITL) is a large-scale data collection project undertaken by CAIDA and DNS-OARC every year since 2006.
  - https://www.dns-oarc.net/ditl/2011/
  - 50 hours packet capture at root DNS servers and other DNS servers (48hours are used by this analysis)
  - Source IP addresses of i.root-servers.net data are anonymized

Year	Start(UTC)	End	Analyzed data from
2011	Apr 12 1200	Apr 14 1200	a,c,d,e,f,h,j,k,l,m (10/13)
2012	Apr 17 1200	Apr 19 1200	a,c,e,f,h,j,k,l,m (9/13)
2013	May28 1200	May30 1200	a,c,d,e,f,h,j,k,l,m (10/13)

JP datasets

- .JP has 1,349,059 registered domain names (on Oct.1, 2013)
- JP DNS servers serve 1.6 billion queries per day
- Two datasets
  - Packet captures of all JP DNS servers, around the same time as DNS-OARC DITL event (and more)
  - Query logs of 2 (a and g) JP DNS servers, every day, for 9 years

Name	Operator	Location	Address (IPv4:7, IPv6:6, total 13)	Capture
A.DNS.JP	JPRS	JP*2	203.119.1.1, 2001:dc4::1	Pcap/Log
B.DNS.JP	JPNIC	JP*1	202.12.30.131, 2001:dc2::1	Pcap
C.DNS.JP	JPRS	Worldwide	156.154.100.5, 2001:502:ad09::5	Pcap
D.DNS.JP	IIJ	JP*2, US*2	210.138.175.244, 2001:240::53	Pcap
E.DNS.JP	WIDE	JP*1,US*1, FR*1	192.50.43.53, 2001:200:c000::35	Pcap
F.DNS.JP	NII	JP*1	150.100.6.8, 2001:2f8:0:100::153	Pcap
G.DNS.JP	JPRS	JP*1	203.119.40.1	Pcap/Log

Analysis method

- Newly developed C program reads pcap files
- It counts number of some kind of queries per each IP address
  - All queries, RD=0 queries, EDNS0 queries,
  - DO queries, name error queries,
  - "." DNSKEY queries (RD=0), “.” NS queries,
  - "." Queries, UDP checksum off queries

Example of temporary output

- I,JP,183.xxx.yyy.zzz,6100,f7bf7ac13fec5289,177,177,1,3,11,0,4,1,0,0,177,177,4096,0
  - It's one of my private validators
  - The address sent com, net, org, arpa, info, gov, biz, jp, cn, ru, uk, de, br, nl, au, pl, tw, kr, ca, eu, fr, … queries to root (64bit variable can hold 64 TLDs)
  - Total 177 queries, RD0 177 in 48 hours
  - non-existent TLD query 1,
  - Root DNSKEY 3, Root NS 11, Root other 0,
  - Signed TLD DS 4, Unsigned TLD DS 1,
  - Non-existent TLD DS 0, Update 0
  - EDNS0 177, DO set 177 (equal to total queries)
  - EDNS0 UDP size 4096
  - UDP checksum off queries 0

Results

Details of 2013 Root Dataset 1

Total pcap entries	29,127,178,041	
IPv4 header checksum error	43,147,205	0.15%
IP length mismatch	188,779	0.00%
UDP	28,117,508,190	96.53%
TCP	138,665,005	0.48%
DNS queries	28,197,821,558	96.81%
Parsed DNS queries (Filtered)	27,771,861,108	95.35%

- TCP ratio: 0.49% TCP/(TCP+UDP)
- Filter
  - – Eliminates the destination address is not the root server's address
  - – Ignore private addresses and link local addresses

Details of 2013 Root Dataset 2

Root	Number of IP addresses	Number of Queries (billion)	Hours
A	5,590,383	3.2	48
C	4,396,938	3.6	48
D	1,259,358	1.0	21.625
E	1,106,155	0.6	45.183
F	3,538,453	2.9	48
H	3,966,149	2.4	48
J	4,191,718	3.8	48
K	4,328,826	3.6	48
L	2,613,861	3.6	48
M	4,354,599	3.0	48
Total	8,547,065	27.8	

Number of IP addrs seen at root 48h

Year	2011		2012		2013	
Data from	10 root		9 root		10 root	
Total	7,591,031	100%	8,989,786	100%	8,547,065	100%
RD0	5,846,612	77.0%	5,859,493	65.2%	6,081,035	71.1%
EDNS0	2,340,543	30.8%	2,906,287	32.3%	3,572,804	41.8%
DO=1	2,018,839	26.6%	2,621,660	29.2%	3,283,728	38.4%
Update	105,131	1.4%	138,778	1.5%	228,633	2.7%
Update Only	71,972	0.9%	99,902	1.1%	179,874	2.1%
Non-existent TLD	2,606,340	34.3%	2,641,072	29.4%	2,619,836	30.7%
Existing TLD	7,361,794	97.0%	8,697,606	96.7%	8,142,126	95.3%
. NS	1,940,015	25.6%	1,871,995	20.8%	2,082,649	24.4%
. Only	26,877	0.4%	36,920	0.4%	105,784	1.2%
. DNSKEY (RD0)	14,092	0.2%	43,782	0.5%	269,390	3.2%
. DNSKEY . Only	571	0.0%	2,828	0.0%	64,612	0.8%

Number of IP addrs seen at root 48h

Year	2011		2012		2013	
Data from	10 root		9 root		10 root	
Total	7,591,031	100%	8,989,786	100%	8,547,065	100%
RD0	5,846,612	77.0%	5,859,493	65.2%	6,081,035	71.1%
EDNS0	2,340,543	30.8%	2,906,287	32.3%	3,572,804	41.8%
DO=1	2,018,839	26.6%	2,621,660	29.2%	3,283,728	38.4%
Update	105,131	1.4%	138,778	1.5%	228,633	2.7%
Update Only	71,972	0.9%	99,902	1.1%	179,874	2.1%

- EDNS0 and DO support is spreading gradually.
  30.8% to 41.8%   and   26.6% to 38.4%
- IP addresses which send UPDATEs are increasing !!!
  1.4% to 2.7%
- Some of them send UPDATE only.
  0.9% to 2.1%

Number of IP addrs seen at root 48h

Year	2011		2012		2013	
Data from	10 root		9 root		10 root	
Total	7,591,031	100%	8,989,786	100%	8,547,065	100%

- About 30% of IP addresses send non-existent TLD queries
- About 24% of IP addresses send . NS (priming) queries
- Probable DNSSEC validators are increasing. 14,092 to 269,390
- However, some of them send "." query only 571 to 64,612

RFC 5011 test ? Configuration only ?

[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Non-existent TLD	2,606,340	34.3%	2,641,072	29.4%	2,619,836	30.7%
Existing TLD	7,361,794	97.0%	8,697,606	96.7%	8,142,126	95.3%
. NS	1,940,015	25.6%	1,871,995	20.8%	2,082,649	24.4%
. Only	26,877	0.4%	36,920	0.4%	105,784	1.2%
. DNSKEY (RD0)	14,092	0.2%	43,782	0.5%	269,390	3.2%
. DNSKEY . Only	571	0.0%	2,828	0.0%	64,612	0.8%

Number of queries at root, 2013, 48h

Kind of queries	IP addresses		Number of Queries	
Total	8,547,065	100%	2.78E+10	100%
RD0	6,081,035	71.15%	2.58E+10	92.79%
EDNS0	3,572,804	41.80%	1.95E+10	70.17%
DO=1	3,283,728	38.42%	1.89E+10	67.90%
Update	228,633	2.67%	7.05E+07	0.25%
Update Only	179,874	2.10%	3.99E+07	0.14%
Non-existent TLD	2,619,836	30.65%	1.17E+10	42.27%
Existing TLD	8,142,126	95.26%	1.52E+10	54.68%
. NS	2,082,649	24.37%	6.47E+08	2.33%
. Only	105,784	1.24%	6.25E+07	0.23%
. DNSKEY (RD0)	269,390	3.15%	8.50E+06	0.03%

Number of queries send from each address, at root, 48 hours

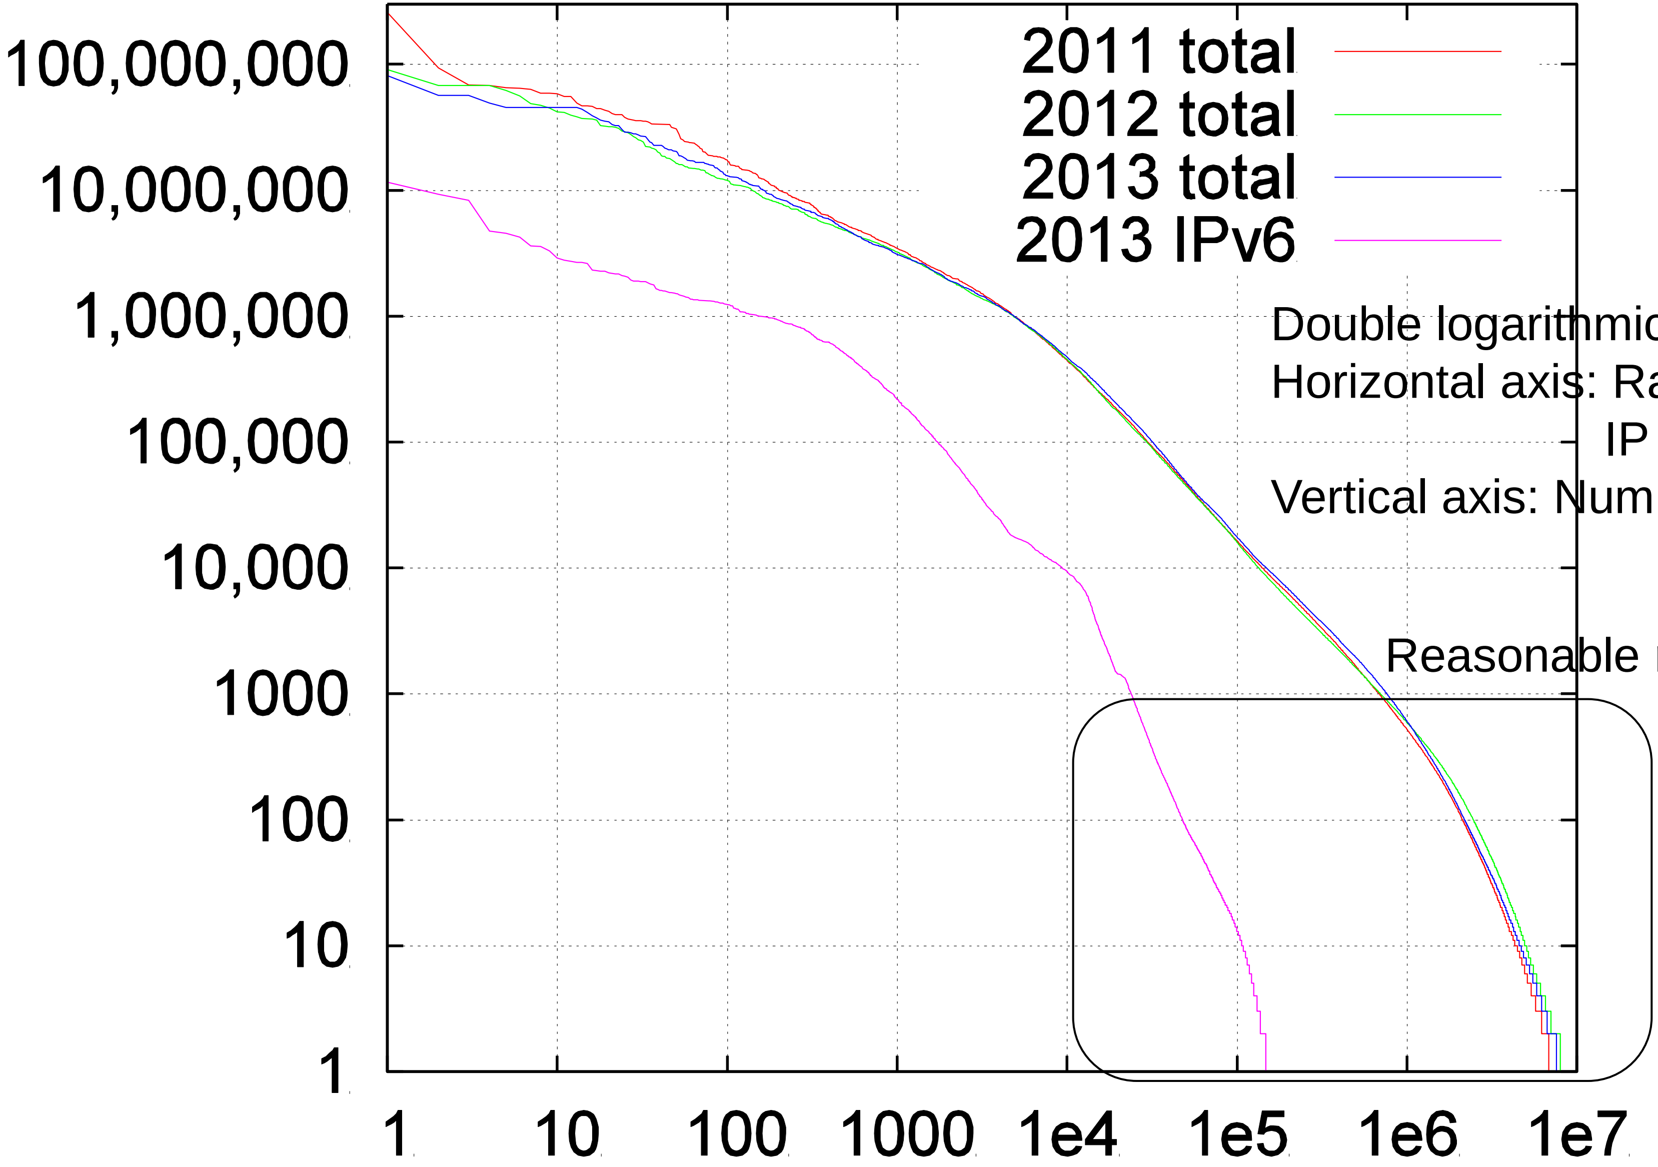

Double logarithmic chart
Horizontal axis: Ranking of IP addresses
Vertical axis: Num. of queries

Reasonable resolvers?

Number of queries from top 50 addresses at root, 2013

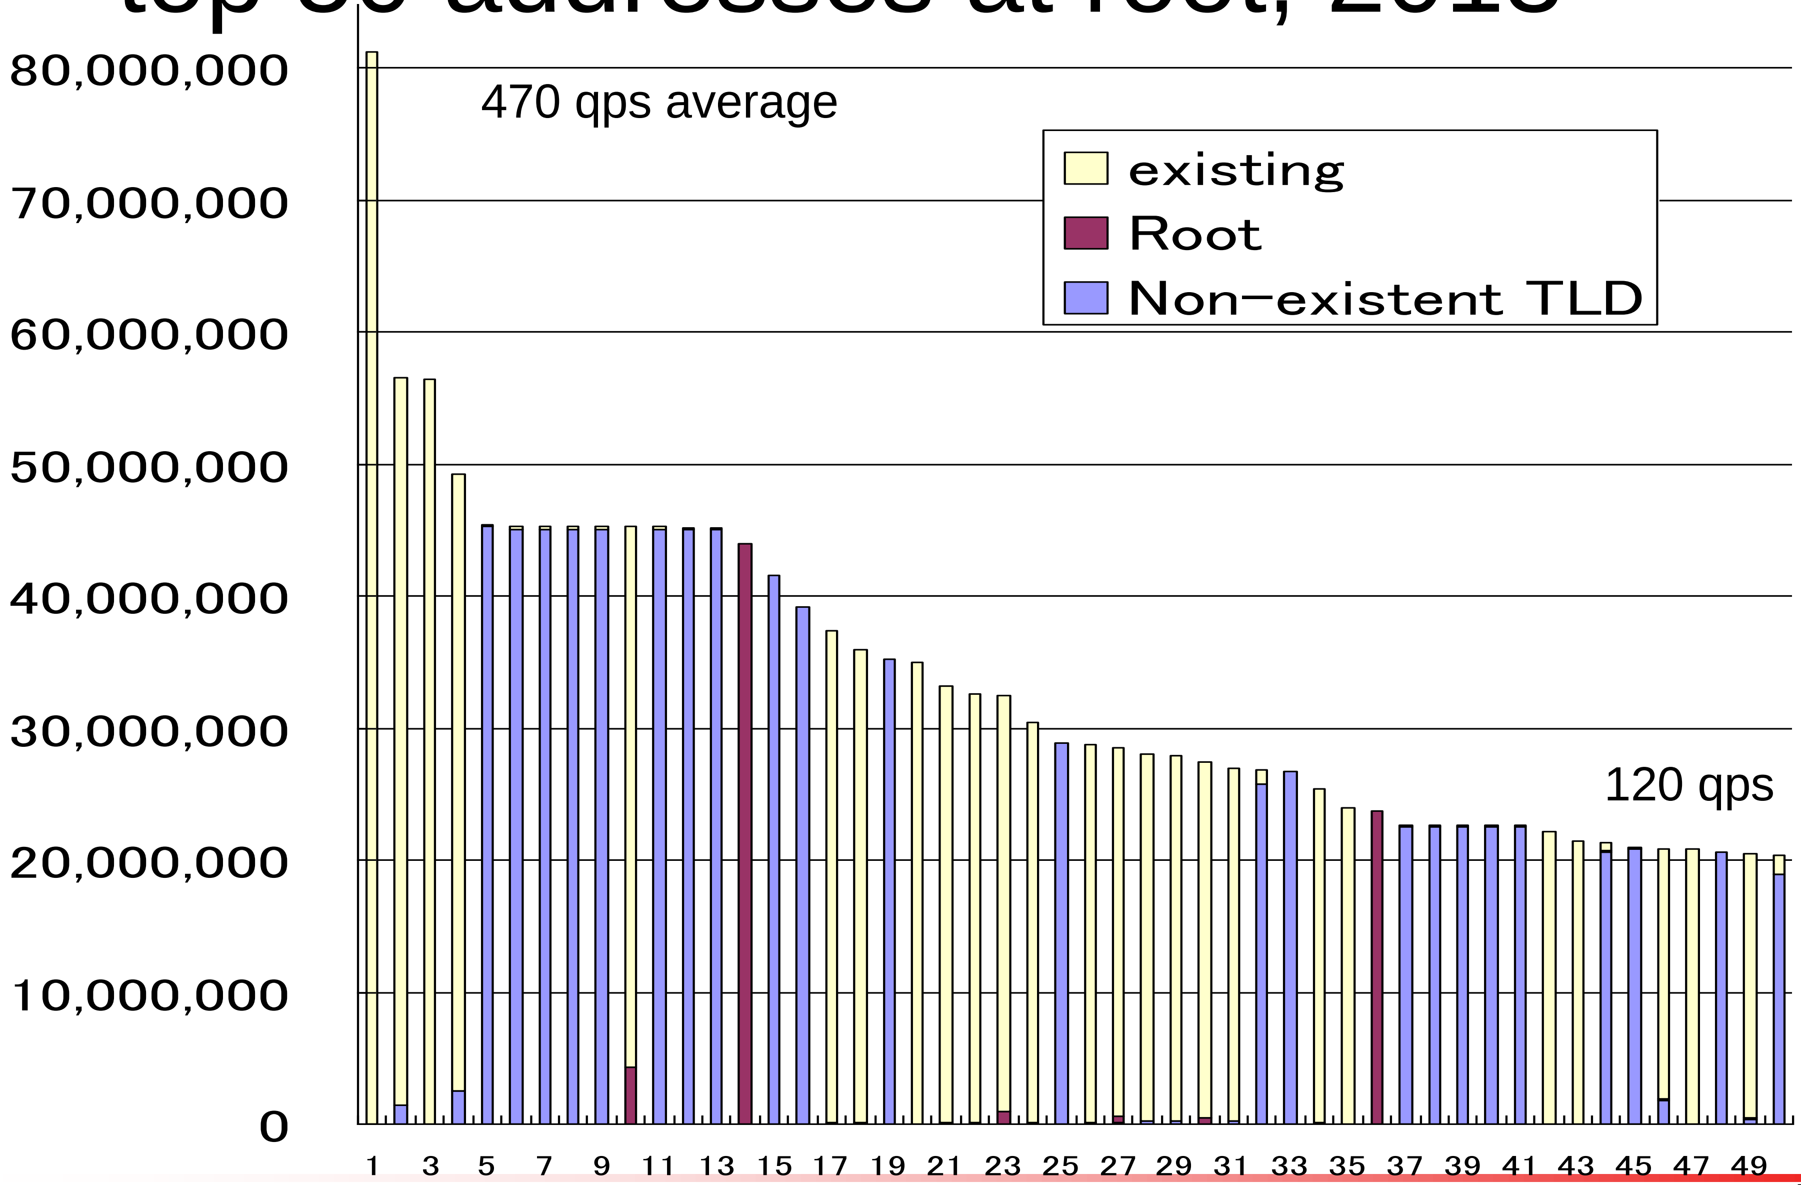

Number of queries from each addresses

- Without TLD typos,
- There are 318 TLDs and their NS TTLs are 172800 and DS TTLs are 86400
  - They should be cached within 1 or 2 days
- If resolvers work well, they should send only 2 * 318 + priming + root dnskey queries at most.
- However, there are 500,000 IP addresses which send over 1000 queries within 48 hours. Why ?
  - They send both existing names and non-existent names

UDP transport analysis

Year	2011		2012		2013	
Data from	10 root		9 root		10 root	
Total	7,591,031	100%	8,989,786	100%	8,547,065	100%
EDNS0	2,340,543	30.8%	2,906,287	32.3%	3,572,804	41.8%
DO=1	2,018,839	26.6%	2,621,660	29.2%	3,283,728	38.4%
UDP checksum off	53,016	0.7%	43,939	0.5%	45,099	0.5%

- In 2013, Number of UDP checksum off queries is 40,368,317, 0.145% of queries.
- Number of IP addresses which disables UDP checksum is small (45,099, 0.5%).
- EDNS0 support is spreading gradually

EDNS0 udp payload size

EDNS0 size	Num of IP Addresses		Queryratio
4096	2,038,920	23.86%	58.73%
4000	900,829	10.54%	4.27%
2048	183,229	2.14%	0.59%
1480	853	0.01%	0.57%
1460	661	0.01%	0.77%
1440	259	0.00%	0.54%
1280	35,816	0.42%	0.31%
1272	990	0.01%	0.66%
1232	630	0.01%	5.29%
1200	4,567	0.05%	0.10%
1024	1,710	0.02%	0.20%
512	78,256	0.92%	1.76%
No EDNS0	5,297,820	61.98%	25.83%

Number of IP address
>= 1000
Or
Queries > 0.5%

Seen at root, 2013

Payload size 4000 and 4096 are used widely.

Another values are seen between
512 and 65535

34.4% of IP addresses support 4000 or more UDP payload size

Comparison of Root and JP 48 hour data

Year	Start (UTC) Day / Hour	Number of unique IP addresses which send Root	. DNSKEY	*.JP query to Root	. DNSKEY and *.JP to Root
2011	Apr12 1200	7,591,031	14,092	1,648,335	7,869
2012	Apr17 1200	8,989,786	43,782	1,346,736	18,858
2013	May281200	8,547,065	269,390	1,502,877	55,659

Year	Start (UTC) Day / Hour	JP	JP DNSKEY
2011	Apr12 1200	2,024,895	5,330
2012	Apr17 1200	1,850,157	14,160
2013	May281200	2,066,612	40,366

DITL 2013 IP address coverage

○ Root (qname *.jp): 1,502,887 (71%)

○ JP: 2,066,612 (98%)

● Intersection(JP∩Root): 1,460,968 (69%)

● Only in JP: 605,644 (29%)

● Only in Root: 41,919 (2%)

Total = Root + Only in JP = 2,108,531 (100%)

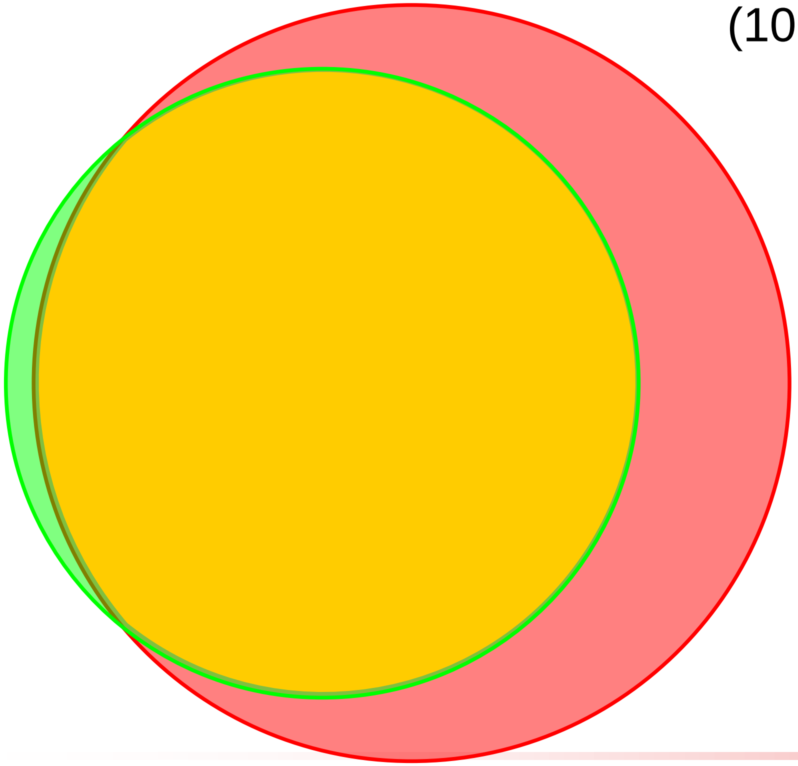

- Compared IP addresses seen at JP and Root, 2013, 48h data,
- JP and Root (qname *.jp) dataset shares 1,460,968 IP addresses
- If JP DNS servers receive most of queries (98% of addresses), DITL-2013 data coverage may be 71% of all IP addresses which interests JP domain name.
- Rest, 29% of IP addresses may be seen at another root DNS servers

Number of queries sent from each address, 48 hours, 2013

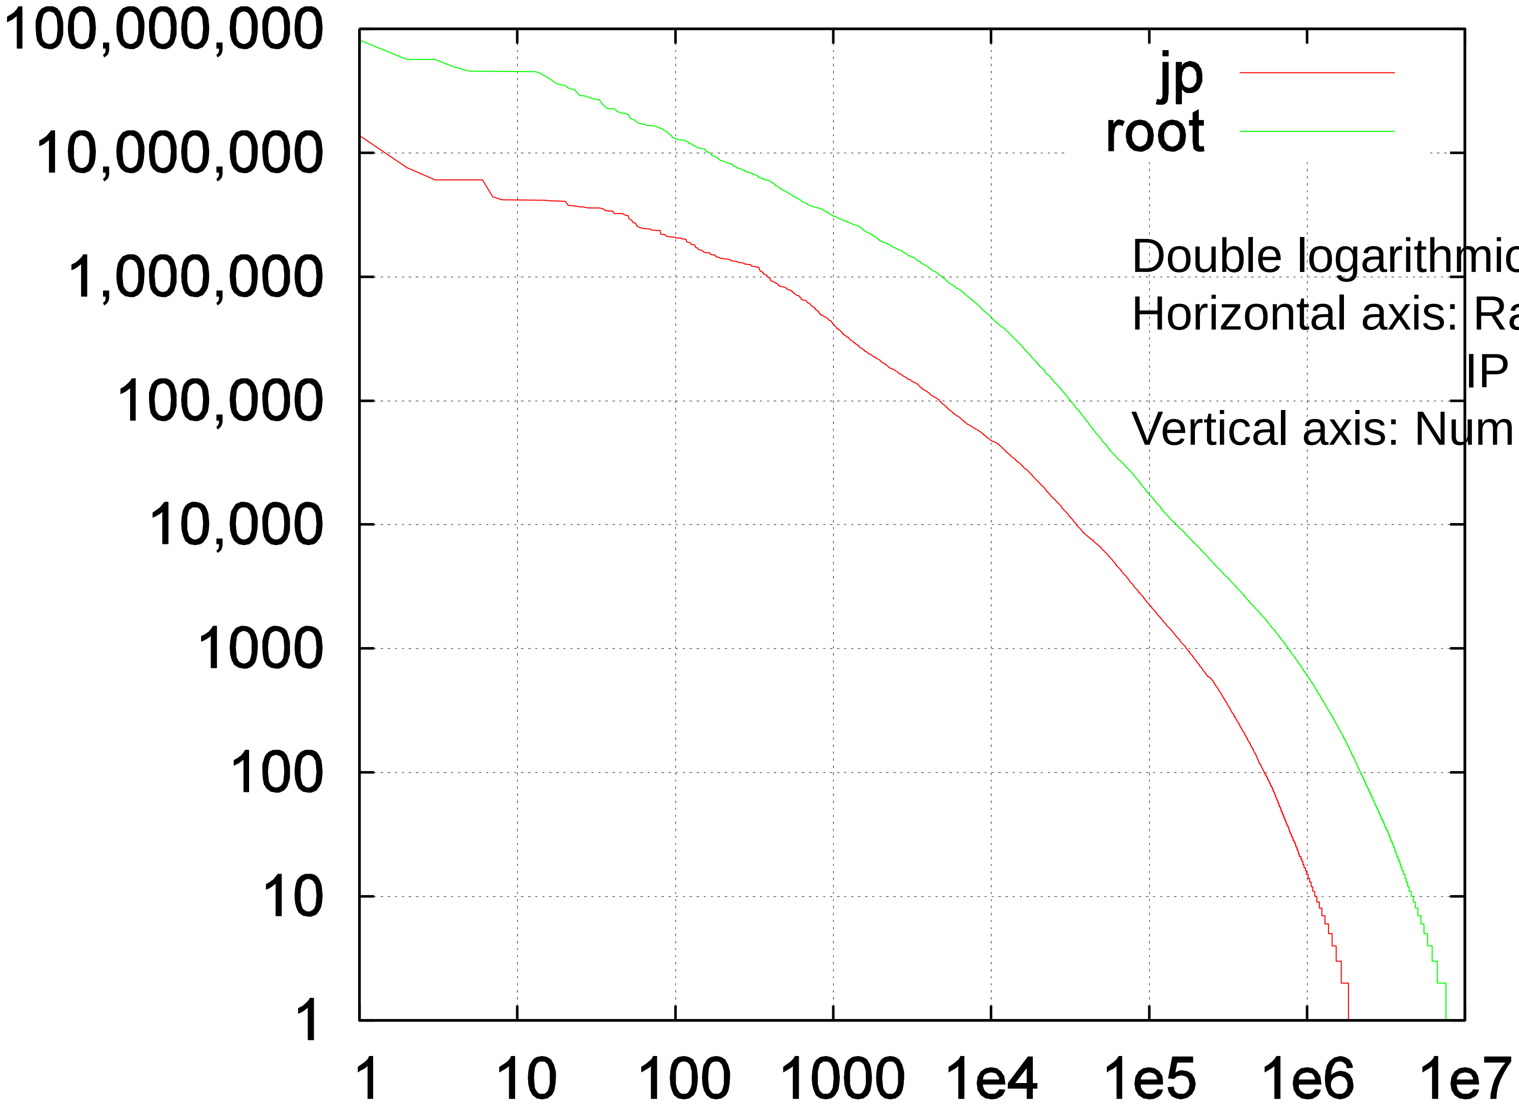

Double logarithmic chart
Horizontal axis: Ranking of IP addresses
Vertical axis: Num. of queries

Number of queries from top 50 addresses at JP, 2013

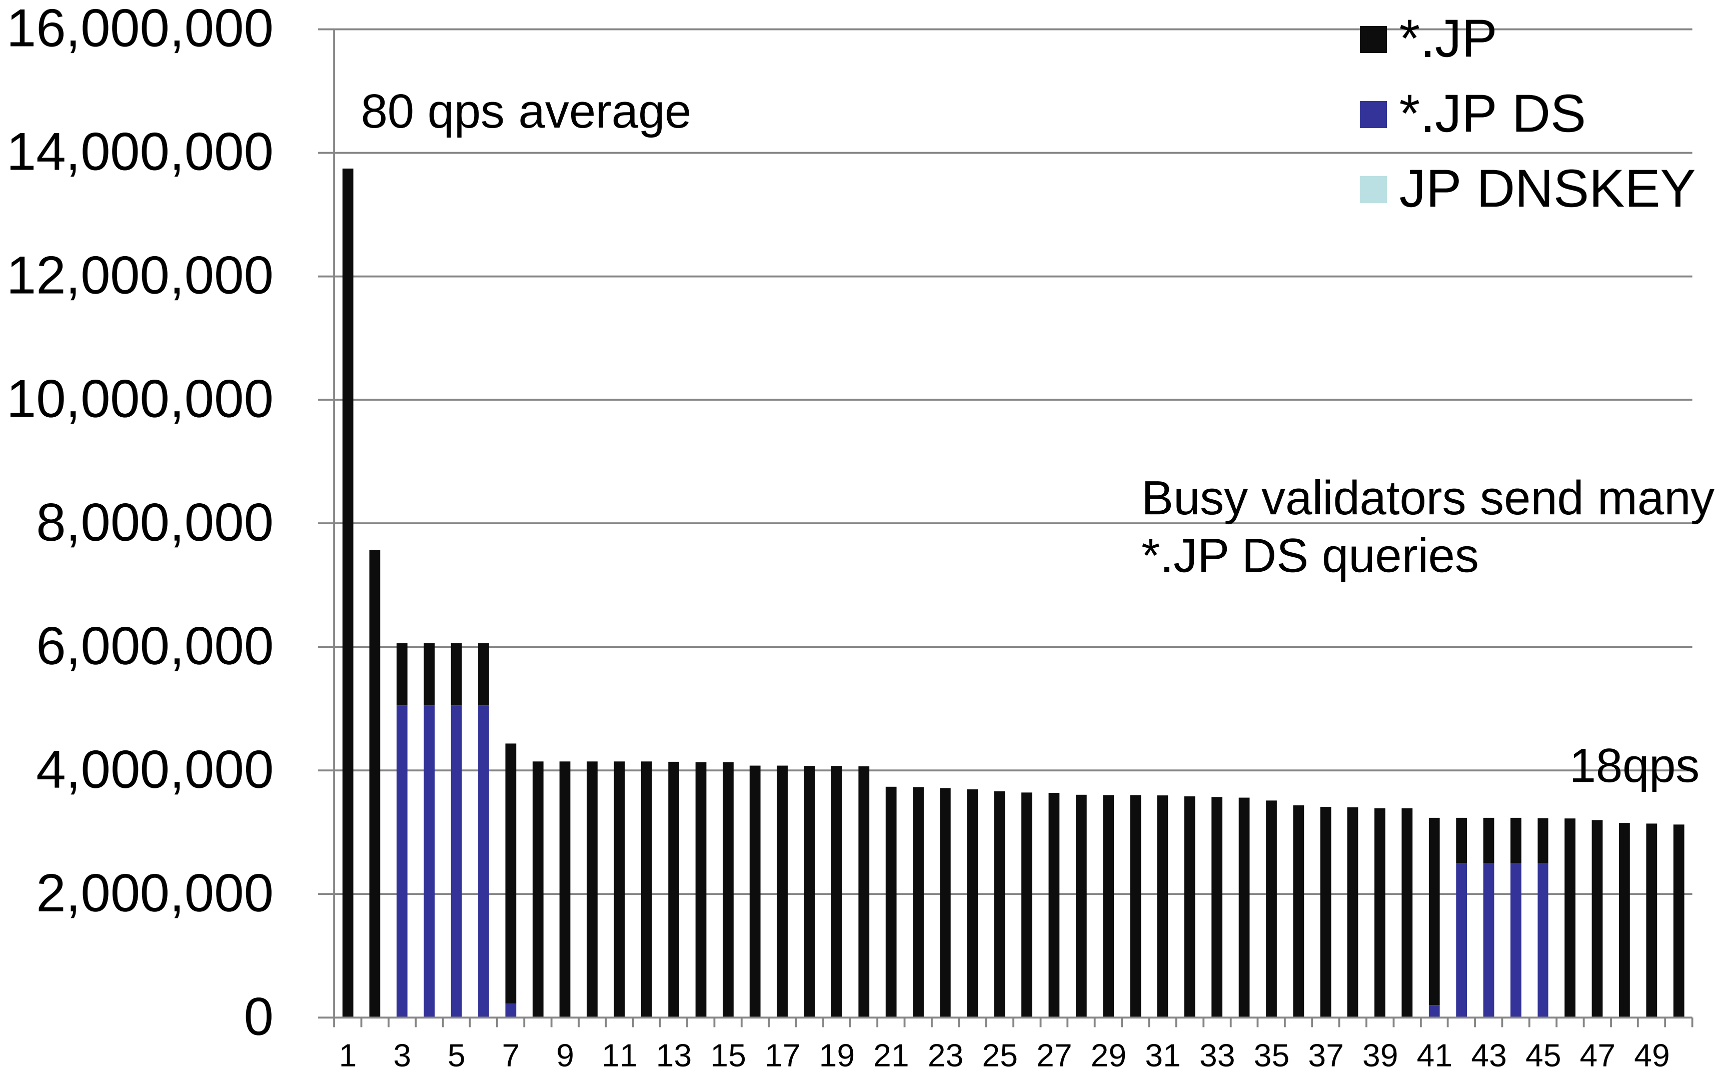

80 qps average

Busy validators send many *.JP DS queries

18qps

Conclusion

- 500,000 IP addresses send 1,000 or more queries to root DNS servers within 48 hours
  - One IP address send 81,261,512 queries
  - I cannot believe number of queries
- IP addresses which disable UDP checksum are very small (45,099)
- 41.8% of addresses support EDNS0 and 34.4% of addresses support 4000 octets payload size
- DITL-2013 data coverage may be 71%

- My analysis is not yet finished and not enough well. I'm continuing my analysis.
  - Qname and rcode analysis, per country, per AS analysis

Acknowledgements

- DNS-OARC as the data source of Root dataset

Differences from my previous presentation

- Differences from IEPG Berlin material
  - http://www.iepg.org/2013-07-ietf87/4%20-%20IEPG-201307-fujiwara-02.pdf
  - Number of IP addresses which send RD=0/1 queries to root
    - Removed port number filter (previously, some source ports are ignored)
  - *.JP query to Root
    - Changed to allow RD=1 queries
  - JP RD0 at JP
    - Changed to allow RD=1 queries