A method for large PCAP file analysis and preliminary results of JP DNS query measurement

Kazunori Fujiwara, JPRS

<fujiwara@jprs.co.jp>

CENTR Tech & RD workshop

June 4, 2012

Abstract

- I wrote PCAP file evaluation program in C and it parses PCAP files fast and with lesser memory usage.

- Recent results are number of possible DNSSEC validators and a preliminary analysis of JP server selection.

Contents

- JPRS' DNS data collection and analysis environment
- Lightweight PCAP analyze program
- Result 1: number of possible DNSSEC validators
- Result 2: classification of JP DNS clients

Overview of JP

- .JP has 1,291,433 registered domain names (June. 1, 2012)
- JP DNS servers serve 1.6 billion queries per day
- Collecting packet captures and query logs

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

JPRS' data sets

1. JPRS collects two days long full capture of DNS packets twice a year.
   - When
     - Once a year: Same timing as DITL (at DNS-OARC)
     - Some events: .JP signed, JP DS in root, IPv6 day
   - Data Format
     - PCAP files separated for each hour and each anycast node
     - Both queries and responses, gzipped
     - Filename contains server name and anycast node information

2. JPRS has been collecting DNS query log from 2 of 7 JP DNS servers for 8 years
   - Not all JP DNS servers
   - Format
     - BIND 9 querylog format, separated for each day and each node, gzipped

Analysis environment

- To protect query data from an outflow
  - Collected data is stored in one machine and analyzed there
  - Only JPRS researchers have access to this machine
  - The machine only exports analyzed statistics
  - Specs:
    - 30TB disk space (Disks sometimes break……..)
    - 24GB memory
    - 4core Xeon * 2
    - Standard server spec. of two years ago

Analysis of query logs

- They are text files
  - If a node receives 4000 queries/sec,
  - Each log file contains 4000 * 86400 = 345,600,000 lines
- Perl (or another lightweight languages) can handle text files well
  - Writing C program which is faster than perl is hard for text processing.
- I don't have good idea to speed up

Analyzing pcap files

- Each pcap file may contain 4000 * 3600 * 2 entries if a node receives 4000 queries/sec.
- Parsing a large PCAP file is hard
  - PacketQ takes 11GB Memory for the data.
    - I don't know how to count multiple number of IP addresses
    - For example, number of IP addresses which send JP DNSKEY, *JP DS and normal JP queries in the same time.
  - Perl is too slow to analyze binary PCAP file.
- Solution
  - I wrote PCAP analyze program in C
    - No special technique
    - Memory usage is controllable, and it is fast.

PcapParseC

- PcapParse.c 931 lines, PcapParse.h 128lines
- It reads PCAP files
  - It supports gzipped file using popen("gzip –cd < pcap.gz |", "r")
  - It calls callback() function for each DNS packets
  - It can parse response packets as well as query packets
- Interface
  - int parse_pcap(char *pcap_file, int callback(struct DNSdata*, int mode), int flags)
  - int callback(struct DNSdata *d, int mode)
  - struct DNSdata contains parsed DNS packet information

PcapParseC

- Users need to write callback (counting) and main (prepare, output) function in C
- Two example applications
  - pcapgetquery: converts PCAP file to BIND 9 format logfile (402 lines)
  - pcapDNSKEY: counts the number of JP DNSKEY queries, the number of DS queries, and the number of JP queries from each IP address (314 lines)
- It is under development: what I want to use.
  - Sorry, no documentation now
  - But easy to build: ./configure; make
- It is available at:
  - http://jprs.co.jp/lab/people/fujiwara/pcapparsec/pcapparsec-0.3.tar.bz2

Parallel evaluation

- Small C program can analyze each PCAP file fast
- It can be performed in parallel
- GNU make handles multiple jobs well
  - "make –j max_jobs" uses CPU cores effectively
- Writing Makefile which analyzes many PCAP files is easy
- We can gather and analyze Intermediate results (from each pcap file) using various programming languages

(example) Processing time

- Input: 50hours pcap files, about 64billion entries, 728 GB, gzipped
  - Original size is over 2TB
- pcapDNSKEY takes 50 minutes for 50 hours data
  - Make –j 12
- Gathering and analyzing takes 25 minutes
  - Written in perl and slow sequential program
  - It outputs one line result

Preliminary result 1

Number of possible DNSSEC validators seen at JP

Number of possible DNSSEC validators seen at JP

- JP DS RR has been introduced in root zone
- JP DNSKEY TTL is 86400, 1 day
- Thus, DNSSEC Validators send JP DNSKEY query once a day if the validators try to perform JP domain name validation everyday.
- Or, BIND 9 Validators seem to send JP sub-domain name DS queries for JP DNS servers.

- I counted the number of IP addresses which send
  - JP DNSKEY queries, *.JP DS queries and JP queries

Result of full packet capture (24hours)

Date	Begin Time UTC	Number of IP addresses			Number of queries	
		JP	DNSKEY	DS	DNSKEY	DS
2010/10/16	15:00	1185367	745	57	2070	1108
2010/10/17	15:00	1523473	879	69	1561	2233
2010/12/10	5:00	1470601	2310	2432	5532	4867319
2010/12/11	5:00	1108265	2083	2296	6234	2335665
2011/4/12	12:00	1560468	3838	5979	27302	7326974
2011/4/13	12:00	1517979	3699	5826	26110	7295136
2011/6/7	11:00	1557000	4673	6925	34744	9990825
2011/6/8	11:00	1493595	4337	6875	38346	9295877
2011/12/13	0:00	1560377	7528	10046	51198	22308672
2011/12/14	0:00	1576341	7388	9998	50358	22602591
2012/4/17	12:00	1284969	10017	15016	45818	25657095
2012/4/18	12:00	1288713	10147	15198	45933	26187764

Each 50 hour data analysis takes about 75min. I can generate this table in one day.
When new 50-hours data set is available, I can get new result in 75 minutes.

Result of full packet capture

- Number of IP addresses which sent JP DNSKEY queries was 10147, Apr. 18, 2012
  - It increased by 6,000 IP addresses in one year
  - It seems to be increasing
- Number of IP addresses which sent DS queries was 15,198, larger than number of IP addresses which sent JP DNSKEY queries
  - I don't know why. Do you know?
  - About 1.5% of JP queries are DS, now
- Number of IP addresses which sent JP queries had not changed for a year.

Number of possible DNSSEC Validators with extrapolated data from query logs of 2 of 7 JP DNS servers

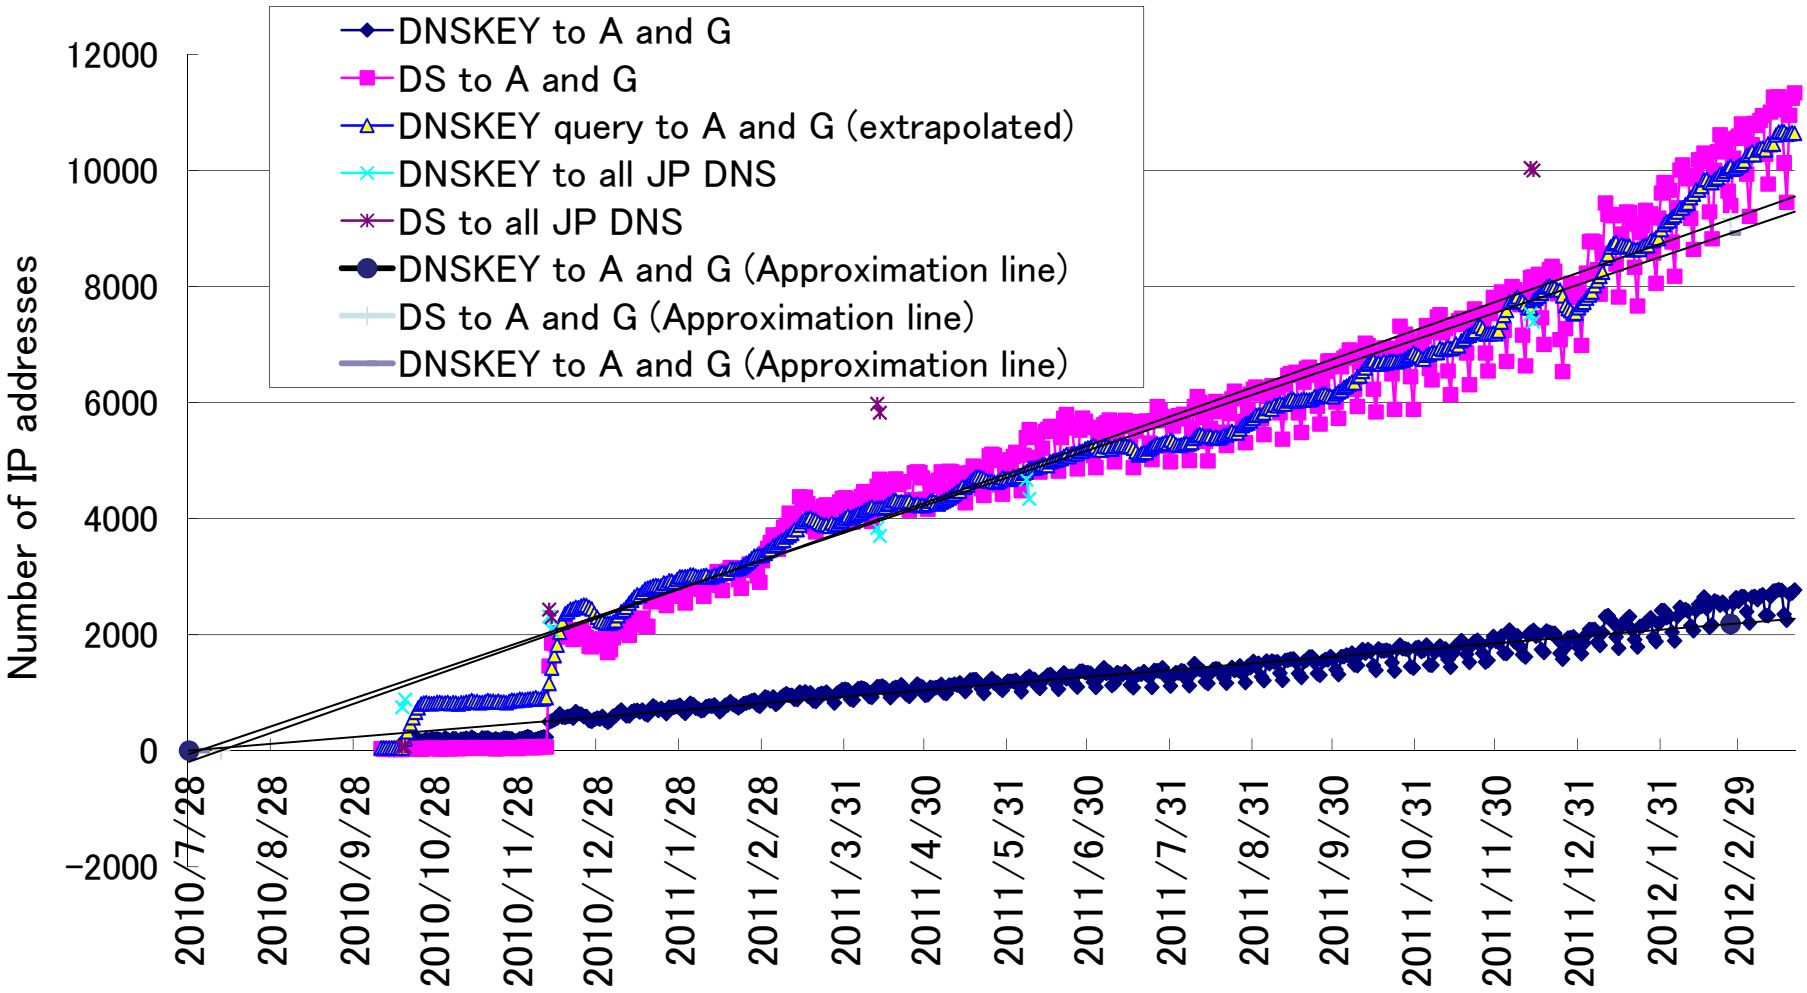

This chart is out of scope of today's main topic because it is result from querylogs.

Result of the analysis

- Details were written in another material
  - See the last IPEG meeting material
  - http://www.iepg.org/2012-03-ietf83/index.html

- The result may be larger than number of real DNSSEC Validators
  - Because there may be many monitors, dig tests, …
  - It shows people's interest

- Then, the result shows the number of DNSSEC Validators, and people's interest about DNSSEC Validation is still increasing linearly or higher.

Preliminary result 2

Analysis of JP DNS server selection

Analysis of JP server selection

- There are seven JP DNS server names
  - 7 IPv4 addresses {A,B,C,D,E,F,G}.DNS.JP
  - 6 IPv6 addresses {A,B,C,D,E,F}.DNS.JP
  - anycast nodes: I did not use this information now.
- Full-resolvers send each query to one of JP DNS servers
- The data used for DNSSEC analysis contain hourly queries from each addresses to each server
- I did two analysis
  - How many queries each full-resolver send for each JP DNS server
  - Which JP DNS servers each full-resolver use for each hour

Cumulative distribution of full-resolver queries

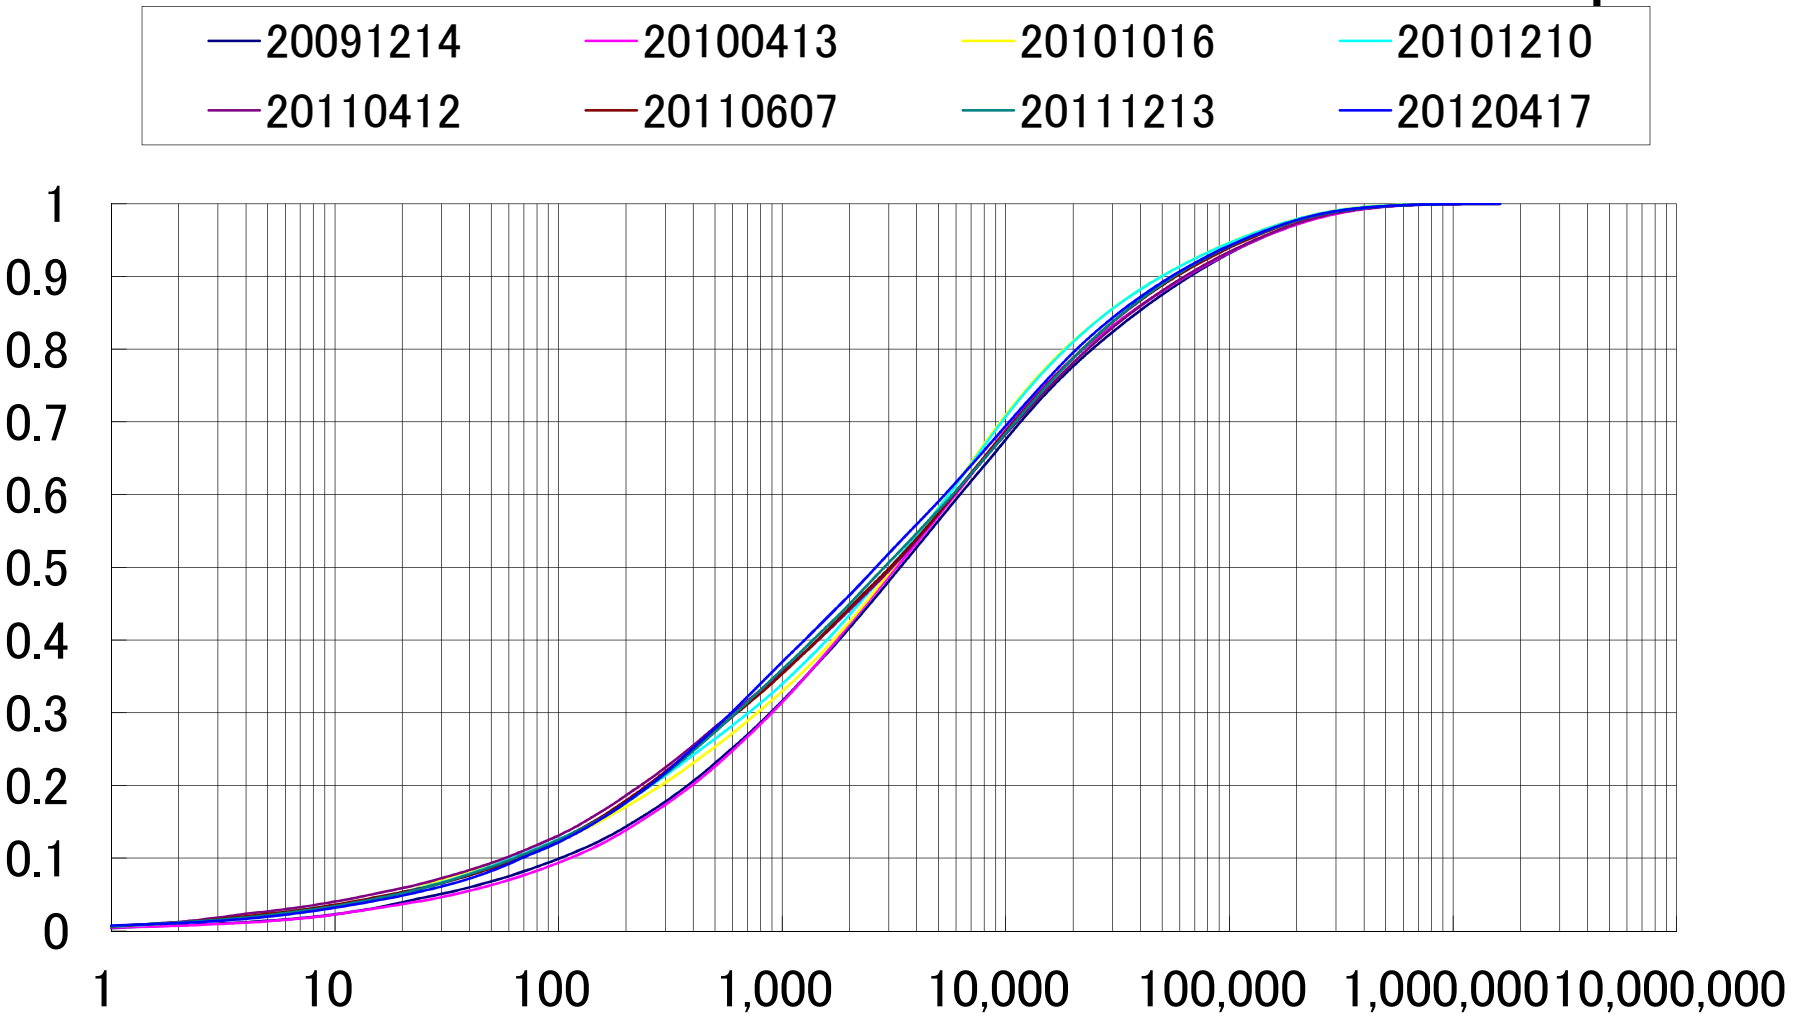

2012/4/17

10,000 IP addresses send 70% of queries, 100,000 IP addresses send 93% of queries.

No 10,000 client generates 47859 queries/50hours, 957 queries/hour

No 100,000 generates 1885 queries, 37 queries/hour

Used JP DNS server in each hour

- 2012/4/17 50-hours data
- From top 100,000 IP addresses (93% of queries)
  - 83,410 addresses use all JP DNS servers in 50 hours, they generates 85.8% of all queries
  - 10,909 addresses use all JP DNS servers in every hours, and they generate 55.8% of all queries
- From top 10,000 IP addresses (70% of queries)
  - 9,090 addresses use all JP DNS servers in 50 hours, and they generates 64.9% of all queries
  - 6,255 addresses use all JP DNS servers in every hours, and they generate 52.1% of all queries
- Most of frequent full-resolvers use all JP DNS servers

Classification of full-resolvers

case	Top 10000 Num. of IP addr.	Top 10000 Queries / All Q.	Top 100000 Num. of IP addr.	Top 100000 Queries / All Q.	Description Each full-resolver sends queries to	Presumption
total	10000	69.45 %	100000	94.23 %		
Equal_7	1013	7%	8463	9.22%	Each JP servers > 0.7 * ave.	They select servers equally
Equal_5	1370	8.99%	12559	12.33 %	Each JP servers > 0.5 * ave.	Or BIND 9.7 TTL BANDING (TTL <128ms)
Japan_7	1402	19.76 %	10171	22.24 %	C < 5% others > 0.7*ave.	Resolvers in Japan? C does not exist in Japan Others are selected equally
C50 %	1878	6.53%	22506	12.33 %	C > 50% of all	Resolvers out of Japan? C is located out of Japan
D50 %	520	4.95%	4376	5.92%	D > 50% of all	Some US resolvers ? D is located at JP and US
C+D60%	3191	14.85 %	36402	23.86 %	C+D > 60% of all	Resolvers out of japan ? C, D (,E) are located out of Japan

2012/4/17-19 50 hours data

Conclusion and Future works of server selection analysis

- Busy full-resolvers use all JP DNS servers
- Some characteristic clusters are easy to find

- Future works
  - Checking the location of each IP address
    - By AS number ? By GeoIP ?
  - Developing good classification algorithm

Conclusion

- PacketQ is useful software, but it is not suitable to parse very large PCAP files
- Writing PCAP file parser written in C is not hard. It works fast with lower (controllable) memory usage
  - http://jprs.co.jp/lab/people/fujiwara/pcapparsec/pcapparsec-0.3.tar.bz2
- Parallel analysis reduces analyze time
  - As a result, over 10TB of PCAP files are able to be analyzed within one working day.