An increase of DS queries to JP DNS servers

and a proposal for its countermeasures

Kazunori Fujiwara, JPRS
<fujiwara@jprs.co.jp>

DNS-OARC 2013 Spring workshop

Contents

- Preconditions
- Overview of JP TLD
- Number of possible DNSSEC validators
- Increase of DS queries
- Reason of DS query increase
- Countermeasures of DS query increase
- Another issue
- Conclusion

Precondition

- "We" are encouraging DNSSEC
  - signing
  - validation
- Most of domains are not signed
- Some of full-resolvers enabled validation

Overview of JP TLD

- JP TLD has 1,334,594 domain names (on May 1, 2013)
- JP zone is signed with 1024bit RSA ZSK/2048bit RSA KSK
- TTL of RRs is 86400, NCACHE TTL is 900
- 7 DNS server names, 7 IPv4, 6 IPv6 addresses
- JP DNS servers serve 1.6 billion queries per day
- Query data collection
  - DITL style packet captures, all servers, 50 hours, 2 or 3/year
  - Query logs on 2 of 7 servers, everyday, 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	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,us,fr	192.50.43.53, 2001:200:c000::35	Pcap
f.dns.jp	NII	jp	150.100.6.8, 2001:2f8:0:100::153	Pcap
g.dns.jp	JPRS	jp	203.119.40.1	Pcap/Log

My past DNSSEC researches

"Counting number of possible DNSSEC validators"

- DNSSec Validation Measures "How to count validators", 13 March 2011, DNS-OARC 2011 San Francisco Workshop, San Francisco, US.
  - IP addresses which send JP DNSKEY queries are possible DNSSEC validators
- Number of possible DNSSEC Validators seen at JP, 1 year difference, 25 March 2012, IEPG Meeting, Prague, Czech
  - IP addresses which send JP domain name DS queries are possible DNSSEC validators

Number of possible DNSSEC Validators

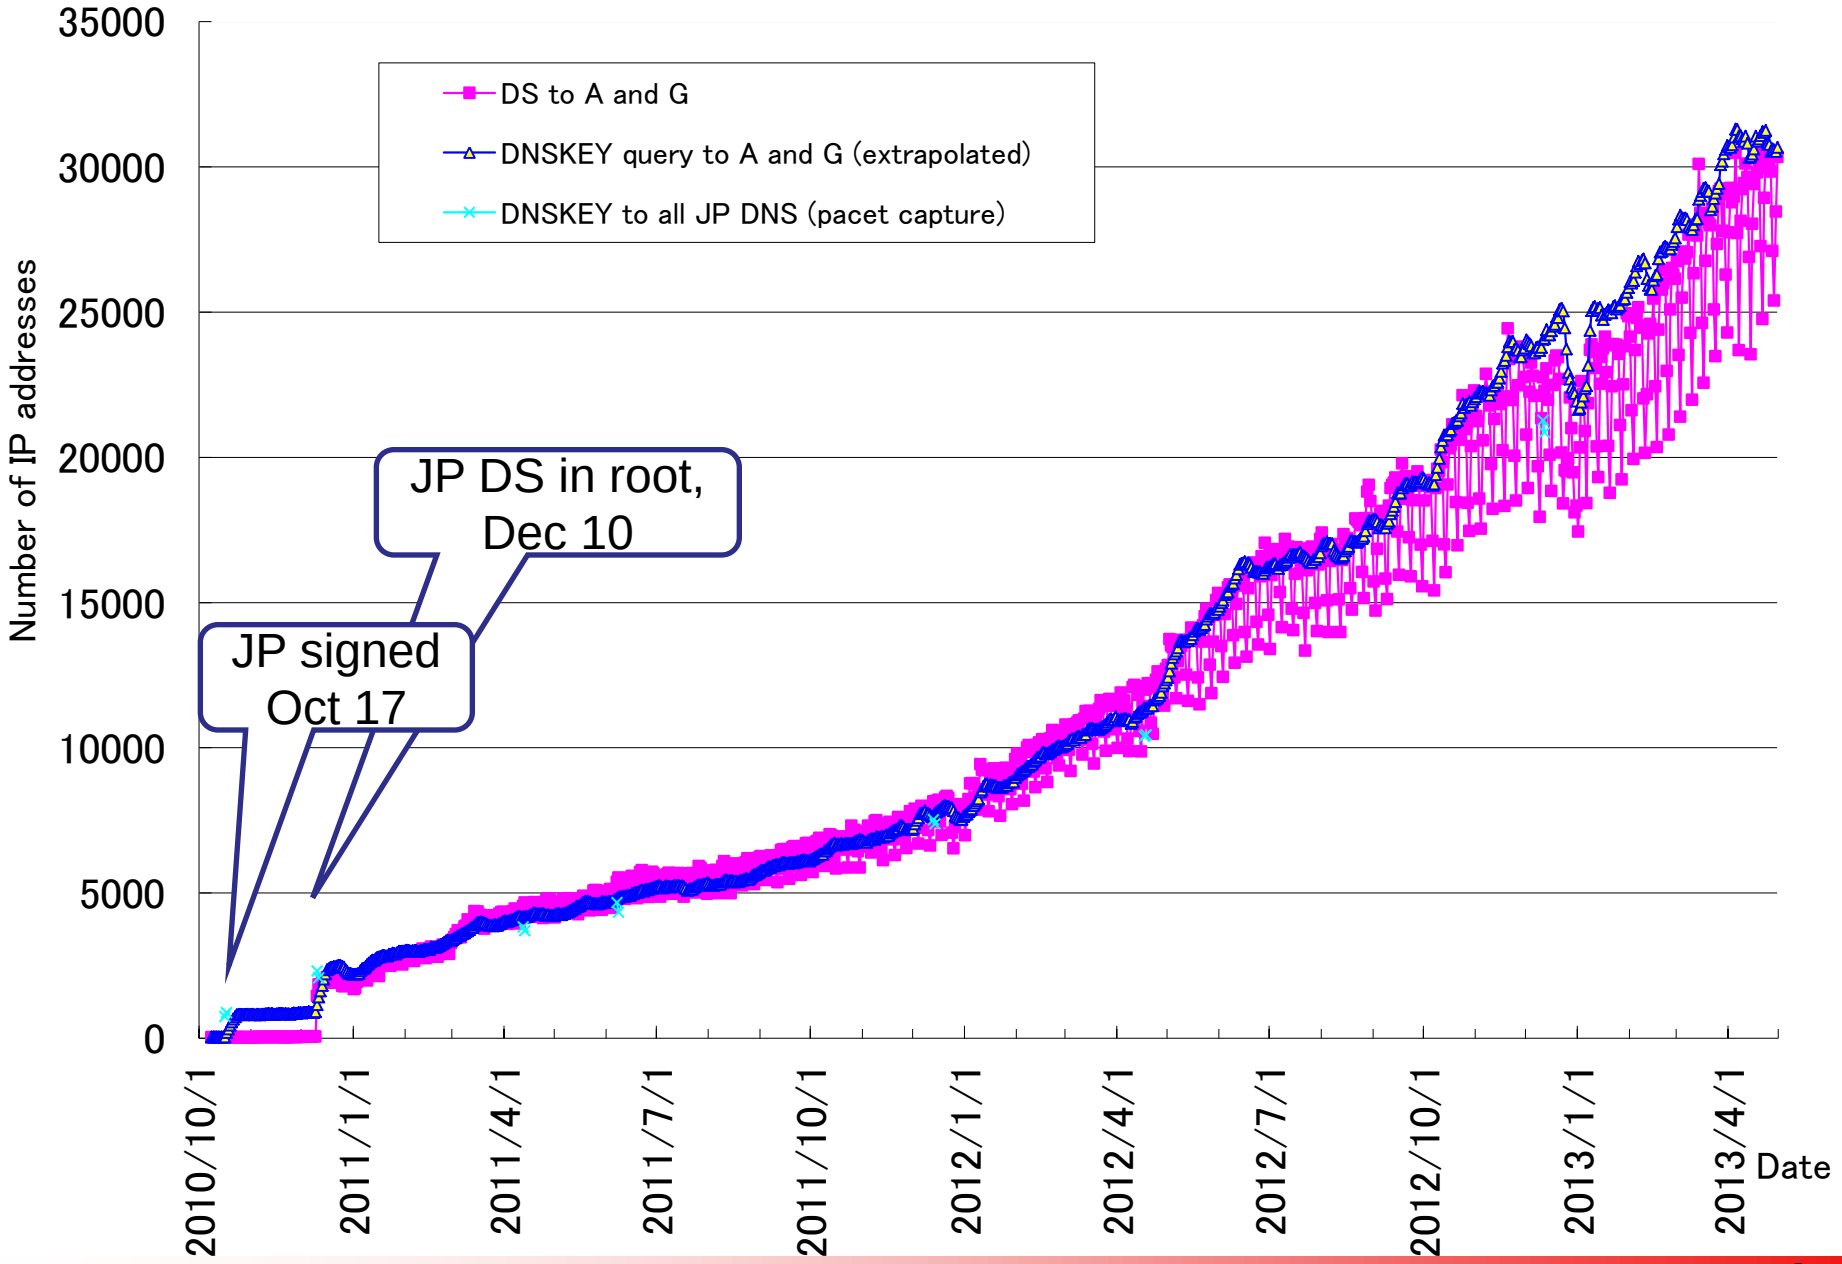

Current ratio of DS queries is 3.5%, Why?

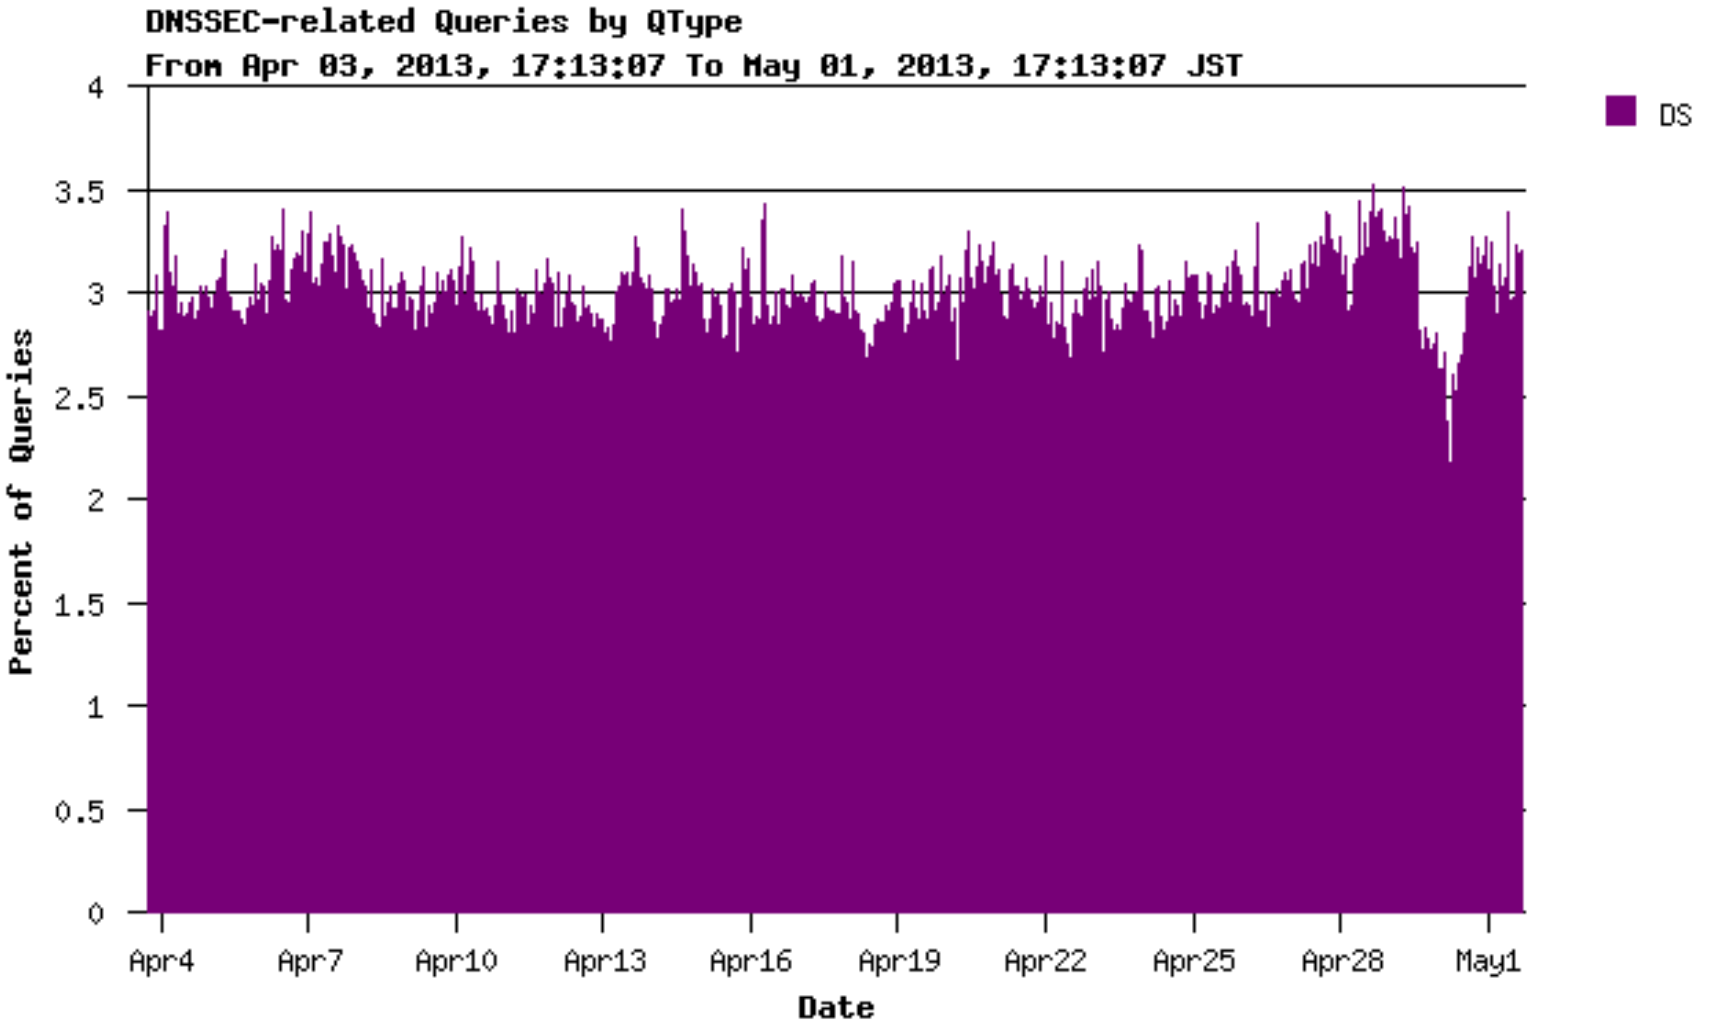

Increase of DS queries (2 of 7 DNS servers)

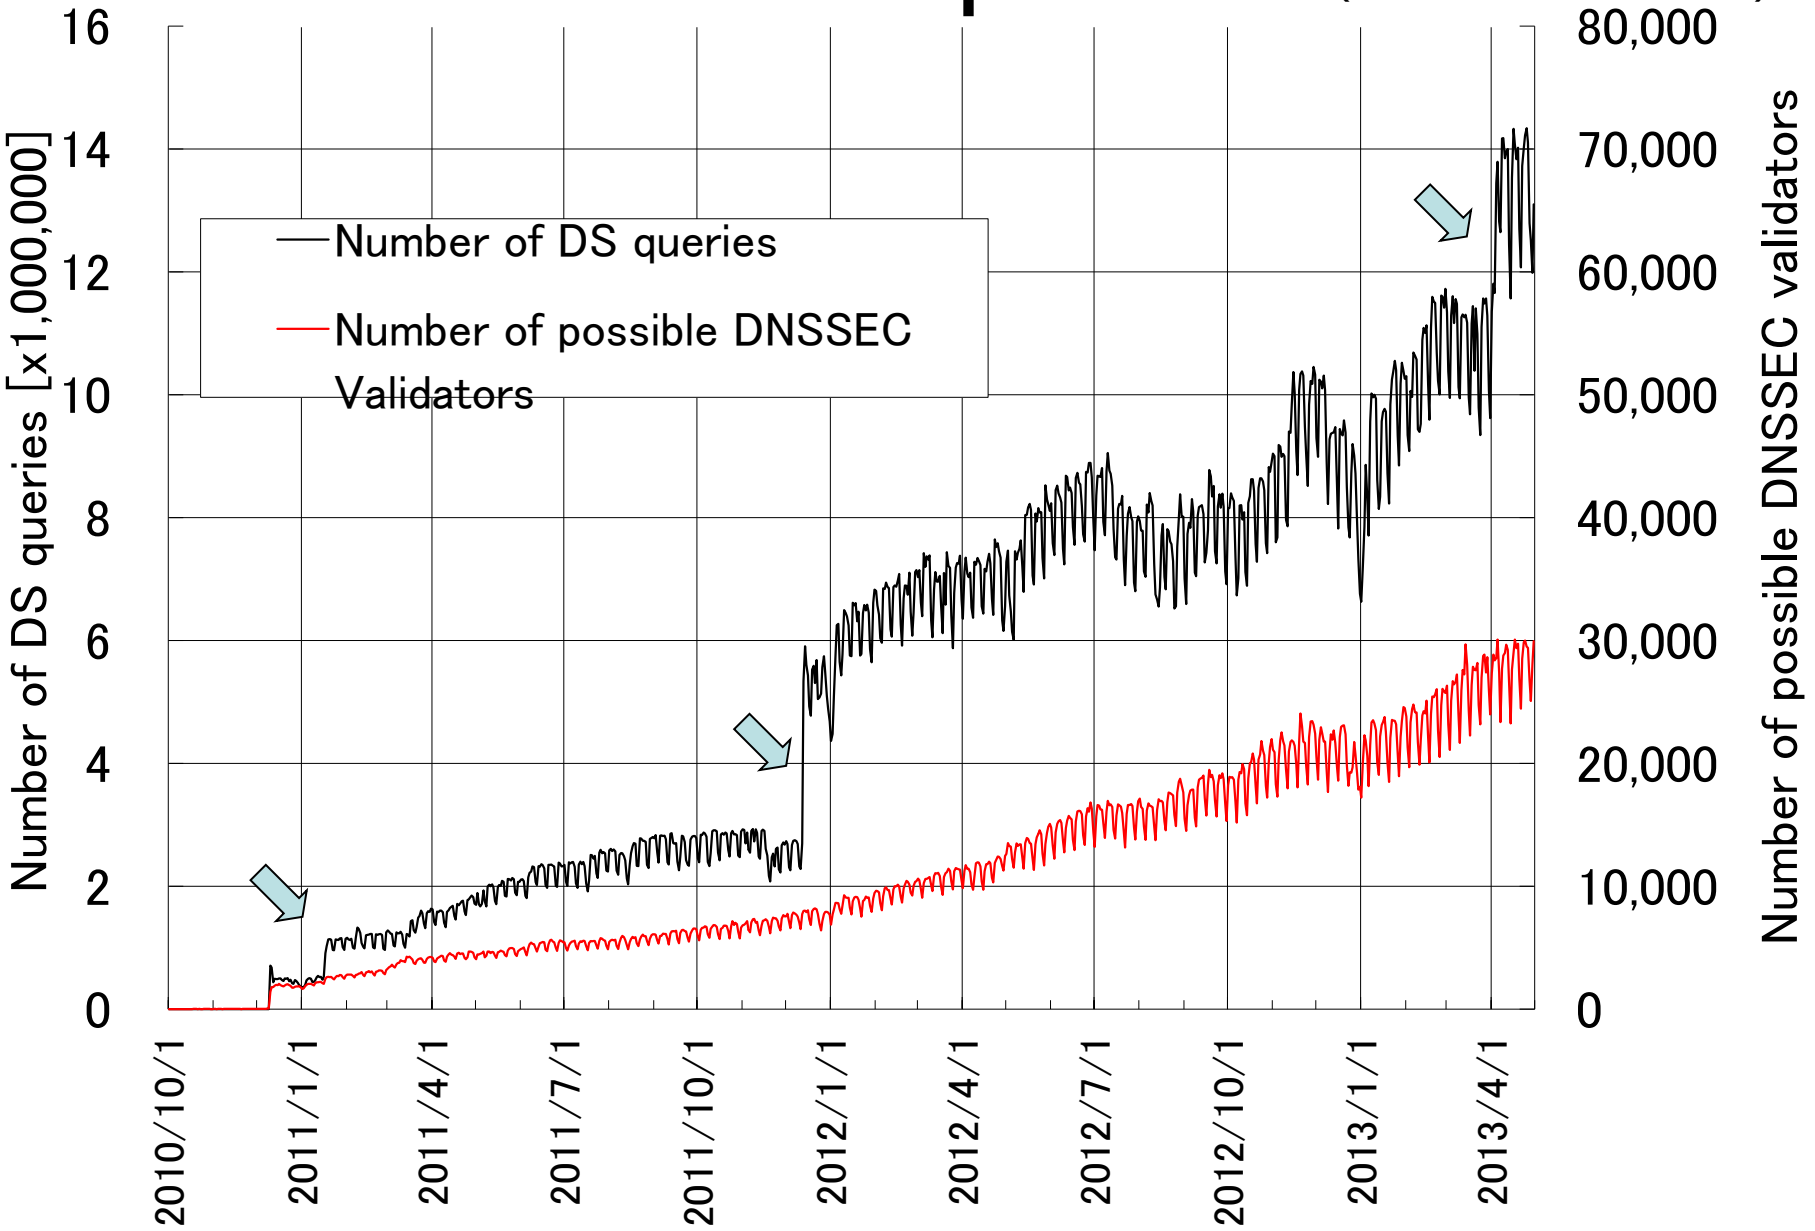

Increase of DS queries (2 of 7 DNS servers)

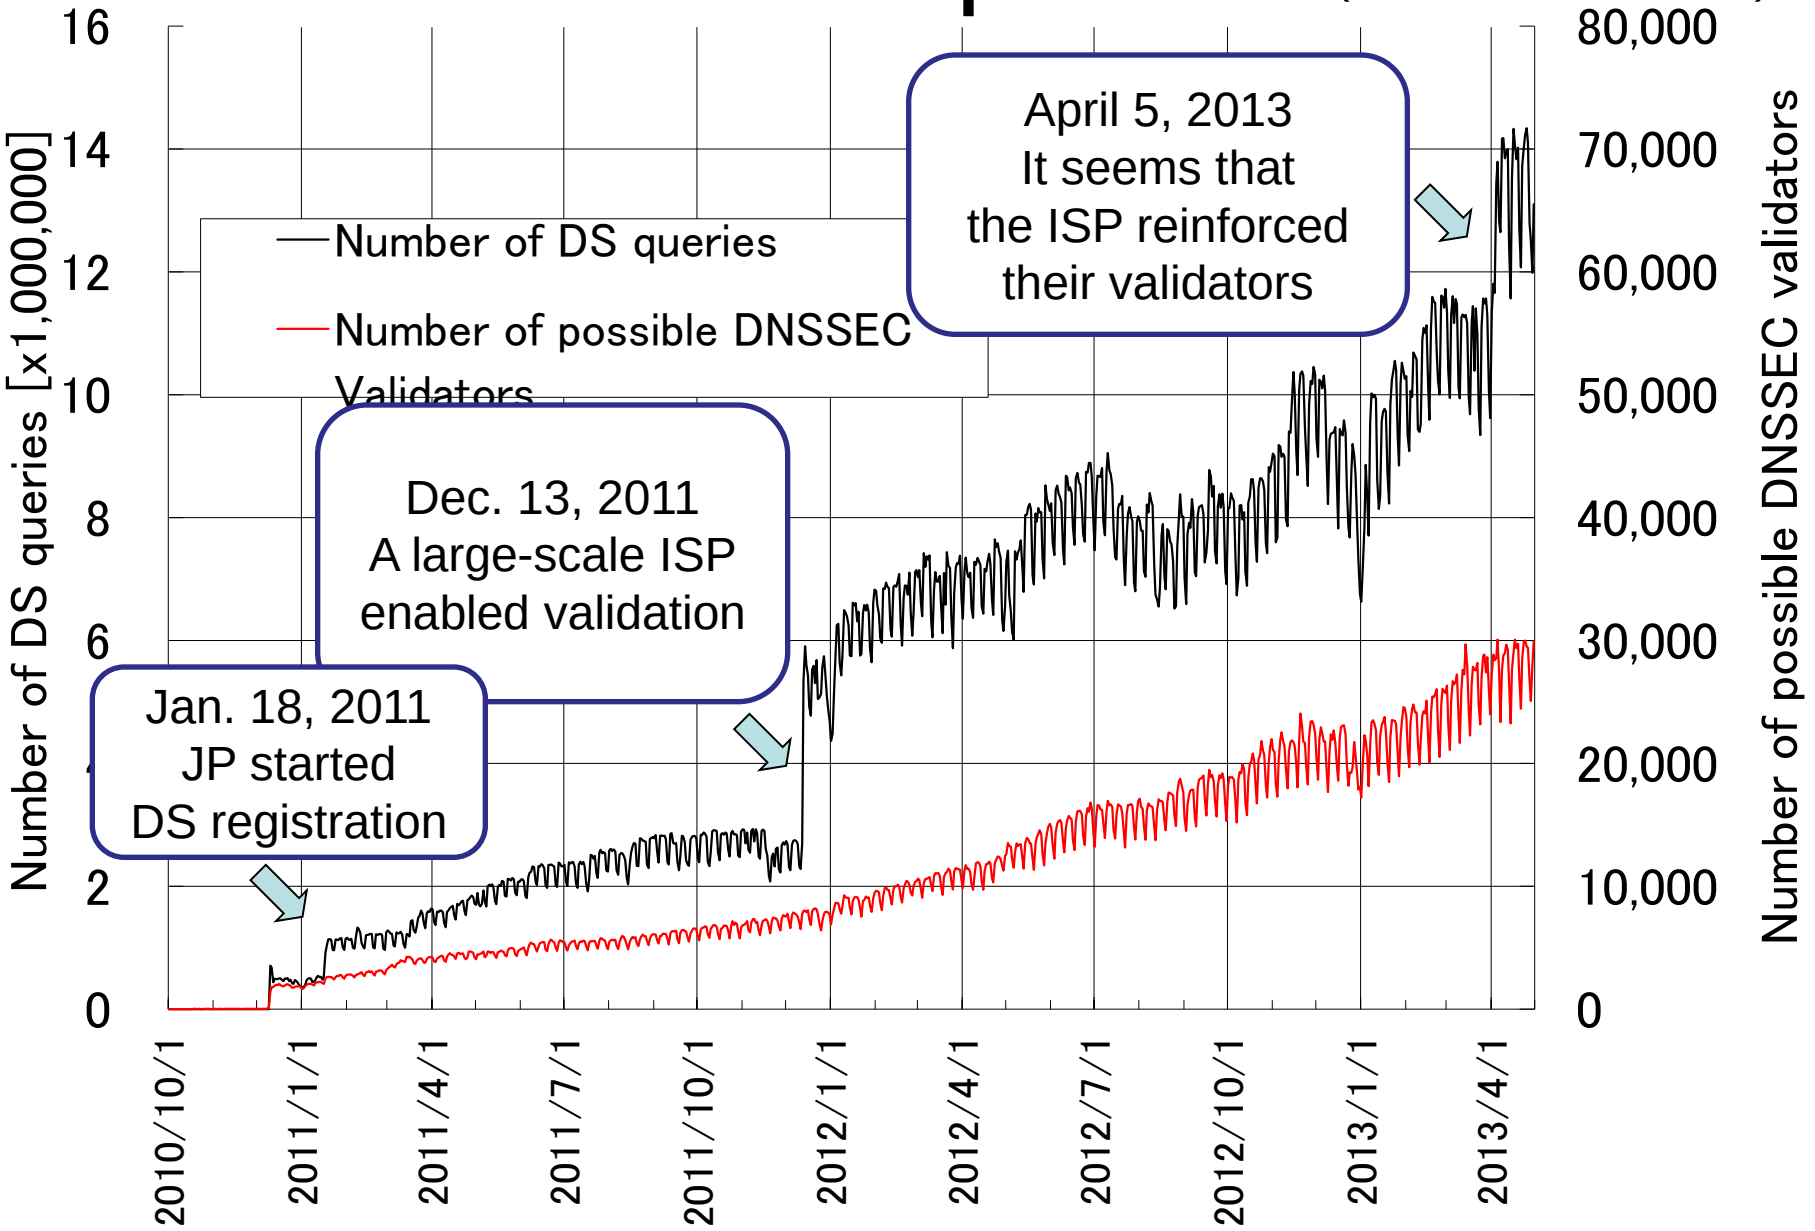

Cumulative distribution of query source IP addresses(2 of 7 JP DNS, Apr 30, 2013)

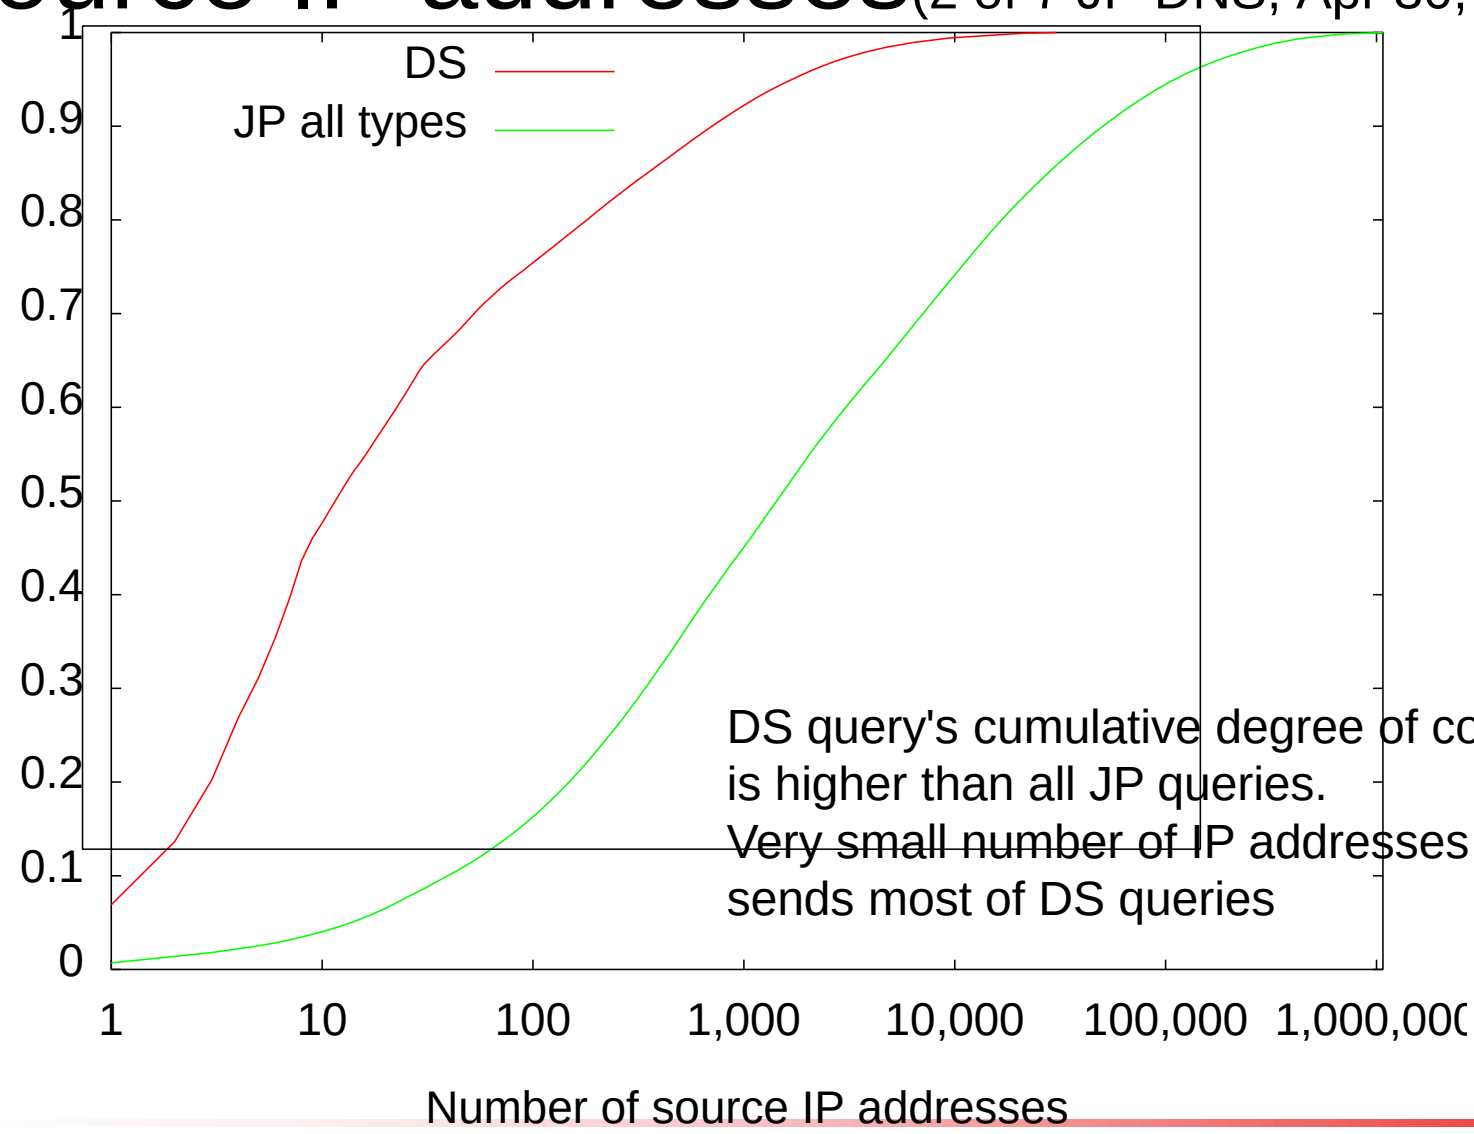

DS query's cumulative degree of concentration is higher than all JP queries.
Very small number of IP addresses sends most of DS queries

A part of query log for a popular name from one IP address, 2 of 7 JP servers

30-Apr-2013 00:19:00.126 google.co.jp IN DS
30-Apr-2013 00:49:00.093 google.co.jp IN DS
30-Apr-2013 01:34:00.369 google.co.jp IN DS
30-Apr-2013 01:49:00.242 google.co.jp IN DS
30-Apr-2013 02:19:01.047 google.co.jp IN DS
30-Apr-2013 02:28:35.867 id.google.co.jp IN AAAA
30-Apr-2013 02:34:01.736 google.co.jp IN DS
30-Apr-2013 03:19:05.265 google.co.jp IN DS
30-Apr-2013 03:34:06.405 google.co.jp IN DS
30-Apr-2013 03:49:08.541 google.co.jp IN DS
30-Apr-2013 04:34:09.628 google.co.jp IN DS
30-Apr-2013 05:04:09.216 google.co.jp IN DS
30-Apr-2013 05:19:09.723 google.co.jp IN DS

One IP address sends many same (google.co.jp) DS queries.
Minimal time interval is 15 minutes, same as JP NCACHE TTL

Reason of DS queries increase

- JP NCACHE TTL is 900, NS, DS, glue TTL is 86400
- Most of JP domain names are not signed
  - They do not have DS RR in JP zone
  - DS RR nonexistence (NSEC3) is cached for only 900 seconds
- Assume there is a popular query name
  - which a validator receives 1 or more queries per NCACHE TTL
  - whose RR TTL is smaller than NCACHE TTL
  - whose domain name is not signed
- Then,
  - Validating process starts every NCACHE TTL period or more
  - The validator need to know DS non-existence every NCACHE TTL period
- As a result, the validator sends one non-DS query per day and 95 (86400/900-1) DS queries per day for the query name
  - It increases queries 96 times
- This is DNSSEC protocol and parameter issue

"www.google.co.jp" case

As a result, JP DNS servers receive google.co.jp DS query every 15 minutes

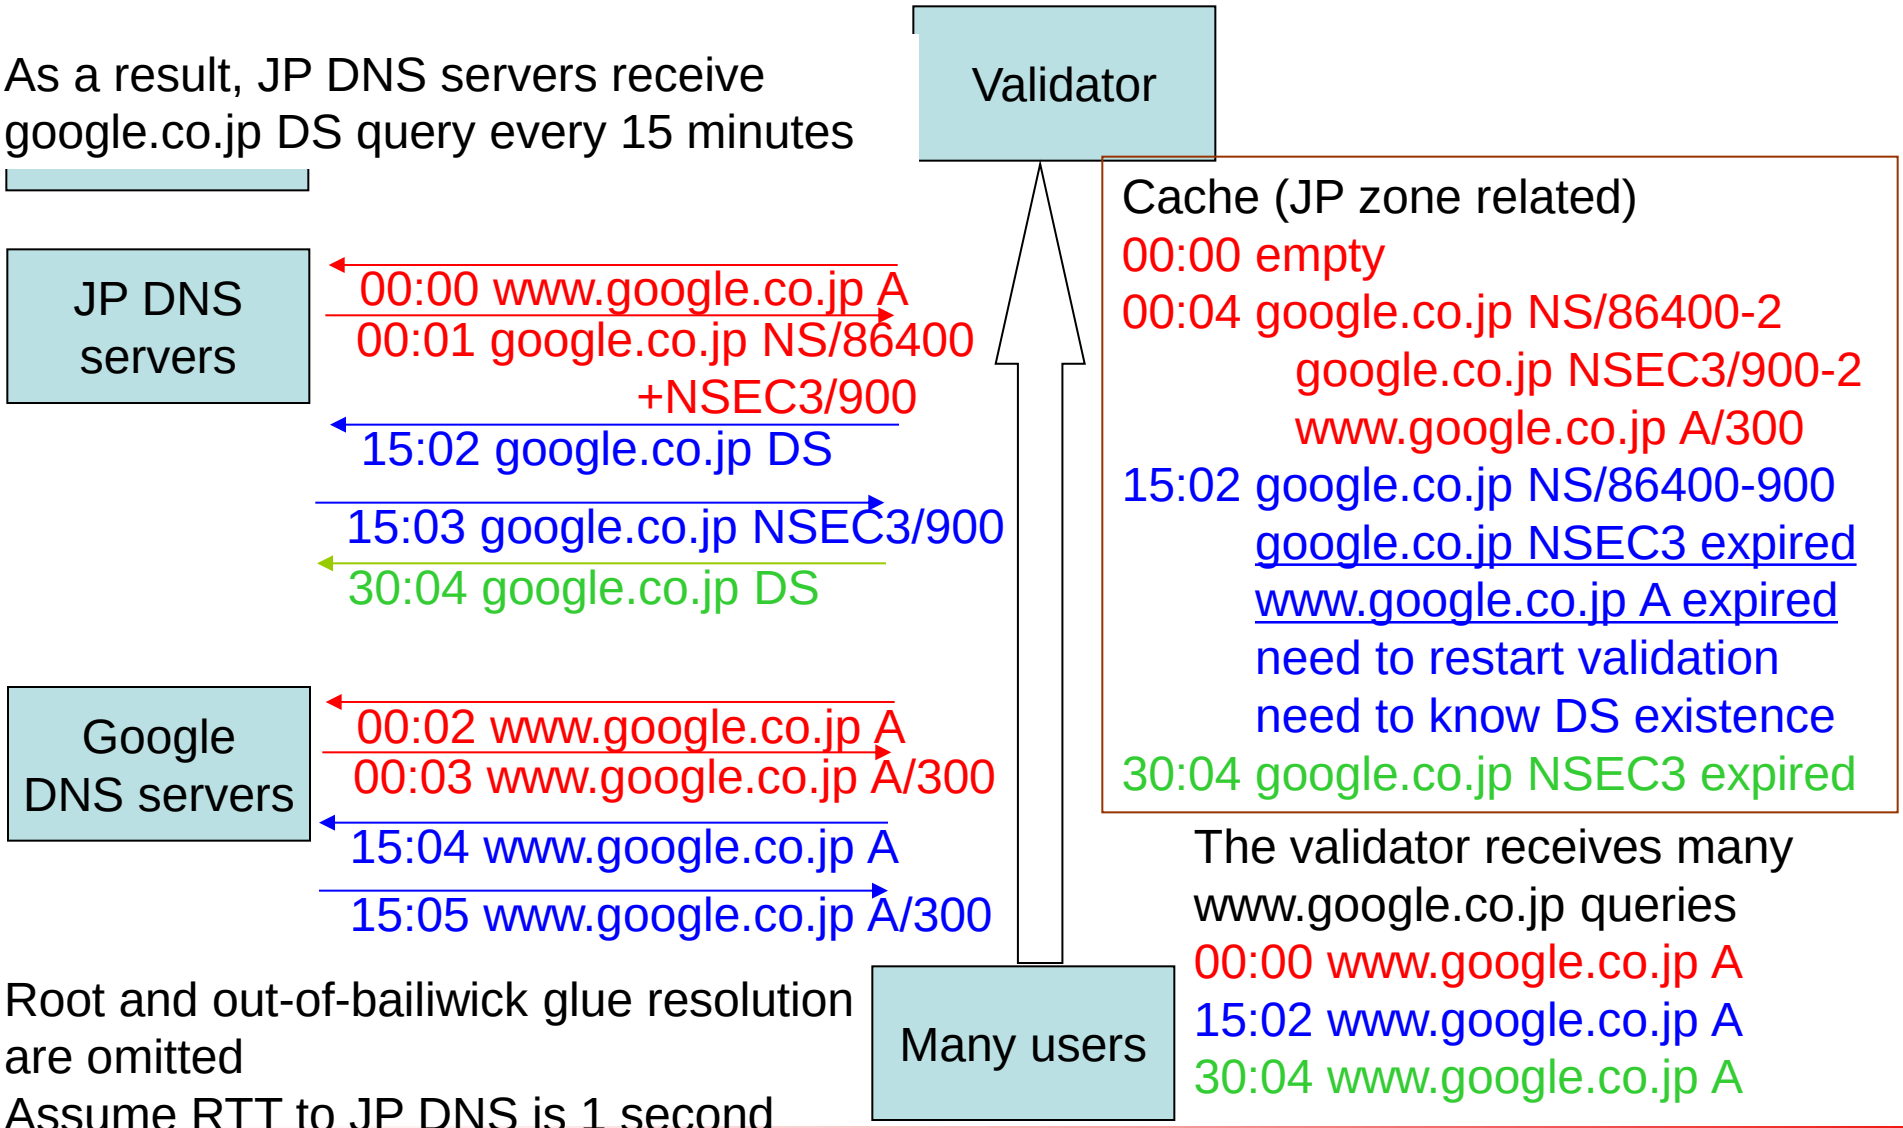

Root and out-of-bailiwick glue resolution are omitted

Assume RTT to JP DNS is 1 second

Impementation Test: BIND 9, Unbound

- Sending periodic queries to test validators
  - dig @validator QNAME A
  - Every 5 minutes
  - Tested QNAMEs:
    - Some of unsigned JP domain names
    - Some of signed JP domain names (jprs.co.jp, jprs.jp)
    - Some of unsigned com, net, org domain names
- Result: Both BIND 9 and Unbound validator send DS queries of unsigned delegation to TLD DNS servers every 15 or 20 minutes
  - Depends on DS existence and RR TTL of QNAME/QTYPE
  - Other queries depend on their TTL

Other observation from logs

- The ISP seems to use multiple DNSSEC validators
  - The validators' caches are not shared
    - Usual usage of full-resolvers / validators
  - 8 IP addresses send DS queries equally
    - They send 40% of DS queries (2 of 7 JP DNS)
    - One IP address sends 1,100,000 DS queries and 250,000 other queries (260,000 unique domain name) to 2 of 7 JP DNS servers
    - It is a common busy full-resolver behavior

Possible situations in the future

- When large-scale ISPs enable DNSSEC validation, their validators start sending periodic DS queries of popular and unsigned delegations

- Therefore, JP DNS servers will receive very large amount of DS queries in the future
  - Magnification is 96 (86400/900)
  - Pessimistically, queries to JP DNS servers will increase 96 times, most of them are DS

Possible affected domains

- Delegation centric zones, signed, small NCACHE TTL
  - RR TTLs of popular names are small (for example, TTL≦600)
- gTLDs
  - Most of gTLDs use 900 as NCACHE TTL
  - TTLs of resource records are 86400 or 172800
  - Magnification factors are 96 or 192
- Some ccTLDs, root, RIRs case (DS TTL / NCACHE)
  - jp: 86400 / 900 96 times
  - co.uk: 3600 / 10800 not affected
  - fr: 172800 / 5400 32 times
  - de: 86400 / 7200 12 times
  - cz: 18000 / 900 20 times
  - us: 7200 / 900 8 times
  - . 86400 / 86400 1 not affected
  - 193.in-addr.arpa 172800 / 3600 48 times

Possible countermeasures

1. Lengthen NCACHE TTL 900 to 86400
   - There is a reason why TLDs chose 900

2. Sign all domain names (so that DS will exist)
   - Possible ? TLD cannot control

3. Lengthen RR TTL of popular names
   - TLD cannot control

4. Add dummy DS to all unsigned delegations
   - Dummy DS TTL value is controllable
   - Need new digest type and deep considerations

Do you have other ideas ?

RFC 4035 Section 5.2

If the validator does not support any of the algorithms listed in an authenticated DS RRset, then the resolver has no supported authentication path leading from the parent to the child. The resolver should treat this case as it would the case of an authenticated NSEC RRset proving that no DS RRset exists, as described above.

Proposal of new digest type

- Define new digest type
- The digest type claims unsigned delegation
- Add dummy DS for all unsigned delegations

- Existing DNSSEC validators do not support newly defined digest type and they should treat the delegations as unsigned
  - I'm afraid that the child zones do not have DNSKEY RRs and validators allow or not (RFC 4035 does not describe well)

Evaluation of dummy DS RR

- A delegation which has dummy DS RR
  - test.dnslab.jp. IN DS 0 0 255 FFFFFFFF
    - Key tag 0, Algorithm 0, Digest type 255
  - test.dnslab.jp. zone is not signed
    - It contains “*.test.dnslab.jp A”
- "www.test.dnslab.jp A" queries
  - Both BIND 9 and Unbound validators resolve well

Side effect of dummy DS RR

1. Zone size will grow
   - JP zone uses NSEC3/OPTOUT
   - If all unsigned delegations have dummy DS RR, all delegations need to be signed
   - Zone sizes will grow, about 10 or more times
     - Signing cost will become very high
     - Zone synchronization will take more time than now
     - Authoritative DNS servers will require more memory
   - However, the change gives same effect as if all delegations are signed
2. DS change (dummy to real) takes TTL time

Things I did not evaluate

- Unknown digest type works or not on other implementations
  - PowerDNS_recursor, Nominum product, etc
  - Other versions of Unbound and BIND 9
  - You can test
    - www.test.dnslab.jp, www.test.co.dnslab.jp
- Is this a general issue ?
  - Do you get many DS queries ?
  - Currently, my validator is sending DS queries periodically to jp, com, net, org DNS servers on my evaluation

Another issue

- Top 4 DS query names (Apr. 30, 2013)
  - co.jp 0.9% empty non-terminal
  - ne.jp 0.6% empty non-terminal
  - yahoo.co.jp 0.45% popular domain name
  - or.jp 0.28% empty non-terminal
- JP DNS servers receive many DS queries for empty non-terminals
  - In JP zone, there are many delegations of co.jp domain names
  - Dummy DS RR idea may not solve empty non-terminal issue (depends on implementations)
  - Zone separation is one of possible solutions

Conclusion

- TLDs which use small NCACHE TTL value and are signed will receive many DS queries of unsigned delegations
- We need to prepare this issue
  - Periodic check of number/ratio of DS queries
  - Considering some countermeasures
    - Write new document to get new digest type
- If you have interests, please comment