A Measurement Study of Open Resolvers and DNS Server Version

Yuuki Takano[†‡]
ytakano@wide.ad.jp

Ruo Ando[†]
ruo@nict.go.jp

Takeshi Takahashi[†]
takeshi_takahashi@nict.go.jp

Satoshi Uda[‡†]
zin@jaist.ac.jp

Tomoya Inoue[‡†]
t-inoue@jaist.ac.jp

[†]National Institute of Information and Communications Technology

[‡]Japan Advanced Institute of Science and Technology

Abstract

DNS is one of the most important infrastructure of the Internet, but it unfortunately suffers from malicious attacks, such as DDoS and cache poisoning. Study and investigation of currently-deployed DNS servers are needed to implement effective and efficient countermeasure. To cope with that, we sent probing requests to the whole IPv4 address space and collected DNS-related information, i.e., DNS server type distribution, DNS server software version distribution and FQDN distribution of DNS server. The measurement result shows that we obtained the addresses of about 30 million DNS servers, about 25 million open resolvers, and about 7 million DNS servers that responded to software version query request. Furthermore, we reversely looked up the DNS servers' addresses to investigate the distribution of domain names. It revealed that there are many open resolvers in spammer-favored domains. We also discuss the relationship between the DNS amplification attack, a type of DDoS attack that abuses open resolvers, DNSSEC, and its countermeasures. DNSSEC significantly increases efficiency of the DNS amplification attack since its records typically amount to tens of thousand bytes.

1 Introduction

DNS [9] is one of the most important infrastructure of the Internet. It provides a name resolving service, with which Internet users can enjoy human-friendly address notation instead of computer-friendly one. Many Internet services depend on DNS. For example, some content delivery network techniques exploit DNS to efficiently deliver contents.

The original specification of DNS was published in 1983 [12], and it has been updated and extended since then. Nevertheless, the protocol is old enough by now to be abused by malicious parties through the methods that were never considered back then. One significant problem is abusing DNS servers to launch DDoS attack. Furthermore, DNS server software such as BIND is often reported their software vulnerabilities.

The DNS amplification attack [10] is a DDoS method, which can cause vast amount of network traffic to victim network or node. It exploits the fact that the ratio of sizes of DNS query and response is quite different: in extreme cases, query and response size are tens of bytes and thousands of bytes, respectively. The DDoS attack is launched by sending packets with spoofed source address that belongs to victim to open resolvers, a type of DNS servers, that bind any address to the UDP socket and accept recursive queries of DNS, and the servers amplify and reflect queries to victim. For example, CloudFrare, which is a service provider, reported that they got 75 GBps DDoS attack by using the DNS amplification attack in 2013 [7]. It is targeted to Spamhaus of non-profit anti-spam organization, and *ANY* query of `ripe.net` is used to attack. To investigate open resolvers, Open Resolver Project [24] discloses its DNS server measurement results on the Internet. Steve Sntorelli also reported of open resolvers and cassified open resolvers into countries [28], but he investigated only limited number of open resolvers.

In this paper, we mainly focus on investigation of open

resolvers in more detail and optionally focus on version distribution of DNS server software. To reveal it, we measure DNS servers on the Internet by probing whole IPv4 address space. Our measurement and analysis results are summarized as follows:

- obtained addresses of about 30 million DNS servers and 25 millions open resolvers
- obtained about 7 million DNS server versions
- revealed DNS server software version distribution
- revealed that 1st-to-3rd level domain distribution of open resolvers by reversely looking up the discovered DNS server addresses
- discovered that there are many open resolvers on spammer-favored domains

2 Methodology of DNS Measurement

In this section, we present our methodology to measure DNS servers on the Internet. Figure 1 shows the architecture of the DNS server measurement system we designed and implemented. It consists of 4 components as follows.

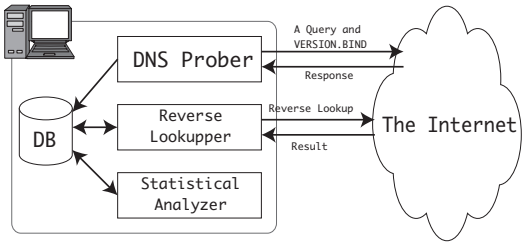

Figure1　DNS Measurement System Architecture

DB

We used MongoDB [18], a NoSQL DB, for our implementation. Measurement results and statistics are stored into the DB. Because of its schema-less feature, we could flexibly develop the DB without being involved with strict record definitions of the DB.

DNS Prober

The DNS prober probes DNS servers in IPv4 address space by sending *A* record requests, whose *RD* flag is unset, to 53 port of UDP. The *RD* flag indicates that querier desires recursive query [9]. If a DNS server receiving *A* query with *RD* flag on for recursive query, it pursues the query recursively and sends the result with *RA* flag, which denotes recursion available, to the querier. Conversely, if a DNS server is unavailable for recursive query, it sends the result of error without *RA* flag.

After receiving a response of *A* query, *TXT* record query of *VERSION.BIND* is sent to the server. Some implementations of DNS server return its software version against it. You can confirm this behavior by running the following *dig* command: `$ dig @127.0.0.1 -t TXT -c CHAOS VERSION.BIND`.

We implemented the DNS prober in C++ using Boost [5], libevent [17], MongoDB C++ Driver and Catenaccio DPI [6]. All the probing results are inserted into MongoDB In our implementation.

Reverse Lookupper

The reverse lookupper reversely looks up IP addresses stored in the DB to obtain fully qualified domain names. We implemented this in C++ by using Boost, libevent, MongoDB C++ Driver and Catenaccio DPI. We took advantage of libevent to query *PTR* records of tens of millions of IP addresses. In our implementation, all of FQDN are also stored into the MongoDB.

Statistical Analyzer

After obtaining results by the DNS prober and the reverse lookupper, the data is statistically analyzed. Calculations for the analysis are performed by MapReduce [16] of MongoDB. MongoDB provides JavaScript language interface for MapReduce, thus we implemented the statistical analyzer in JavaScript.

We probed DNS servers on IPv4 address space by our implementation from 5th to 6th in July 2013. Our measurement revealed that there were about 30 million DNS servers on IPv4 address space, about 25 million of which are open resolvers while 7 millions of which can tell their server software version. More details are discussed in the following sections.

3 DNS Type Distribution

This section describes the measurement results on DNS servers and open resolvers. We classified IPv4 addresses based on regional Internet registry (RIR) [27], and returned strings of *VERSION.BIND* query as DNS types by regular expressions shown in table 1.

Table 2 shows DNS type distribution. Each row denotes DNS type distribution classified as RIR, and each column denotes DNS type distribution classified by table 1. In this table, the “can’t detect” column indicates the number of servers, which rightly returned response against

Table1　Regexs for DNS Type Classification

Type of DNS	Regex
BIND 9.x	^9(\.[0-9])+
BIND 8.x	^8(\.[0-9])+
BIND 4.x	^4(\.[0-9])+
Dnsmasq	^dnsmasq
Nominum Vantio	^Nominum Vantio
Nominum ANS	^Nominum ANS
PowerDNS	^PowerDNS
Unbound	^unbound
NSD	^NSD
Windows series	.*Windows

VERSION.BIND query but the response string from them couldn't be classified by the regular expressions, and the "no version info" column indicates the number of servers, which returned an error message against *VERSION.BIND* query.

We first discuss DNS server and open resolver distribution for each RIR. We obtained 30,285,322 DNS server addresses by sending *A* record queries to whole IPv4 address space, 24,971,990 of which returned responses with *RA* flag. This indicates that about 82.5 % of the DNS servers are open resolvers. Especially, DNS servers of ARIN and RIPE NCC account for about 62.3 % of all the DNS servers, more than 87 % of which are open resolvers. The numbers of DNS servers of LACNIC and AFRINIC are 5,149,451 and 1,205,748, respectively, and more than 96 % addresses of them, i.e. almost all of them, are open resolvers. In ARIN, there are 3,139,392 DNS servers and only 1,720,185 are open resolvers in ARIN; the percentage of open resolvers is less than the others RIR.

We then discuss DNS server types. We obtained 15,357,412 addresses that responded with *VERSION.BIND* query, and 7,075,527 of which were classified by the regular expressions shown in table 1 because response text of *VERSION.BIND* can be modified and configured by operator. BIND series [2], Nominum ANS [20], PowerDNS [25] and NSD [22] are authoritative DNS servers. Table 2 reveals that almost all of the PowerDNS servers and 43.4 % of BIND 9.x servers are open resolvers, but there are few open resolvers of Nominum ANS and NSD. Dnsmasq [11], Nominum Vantio [21] and Unbound [30] aren't authoritative DNS servers; they only work as a caching, resolving or forwarding server. The table reveals that almost all of Dnsmasq and Nominum Vantio are open resolvers, but only 32.3 % of Unbound servers are open resolvers.

We next discuss obsoleted BIND series. The table reveals that BIND 4.x and BIND 8.x series are still alive on the Internet despite the Internet systems consortium, which is the developer of BIND, announced that BIND 8.x series were entering the end of life in August 2007 [3]. It is also revealed that RIPE NCC is the worst holder of obsoleted BIND series. There are 3,486 and 35,218 addresses of BIND 4.x and BIND 8.x, and 2,751 (78.9 %) and 21,348 (60.6 %) of them are in RIPE NCC, respectively. These results imply that software once widely deployed cannot be completely replaced to newer version.

4　DNS Server Software Version Distribution

In this section, we show version distribution of each DNS server types.

4.1　BIND Series

BIND is the most popular authoritative DNS server software. We found 417, 86 and 71 software versions of BIND 9.x, 8.x and 4.x series in total. Figures 2, 3 and 4 show detailed software version distribution for each of the series. We separately counted up pure BIND version and RedHat's one because RedHat independently backports and distributes it for their Linux distribution.

The latest versions of BIND 9.x series are 9.9.3-P2, 9.8.5-P2, 9.7.7 (EOL), and 9.6-ESV-R9-P1 in July 2013 [4]. Figure 2 reveals that many servers aren't updated to the newest versions. The latest versions of BIND 8.x and BIND 4.x are 8.4.7 and 4.9.11, respectively.

4.2　PowerDNS

PowerDNS is DNS server software, and we found 22 versions of it in total, Its software version distribution is shown in figure 5. PowerDNS is distributed as an authoritative server called "PowerDNS Authoritative Server" or resolving name server called "PowerDNS Recursor". Figure 5 shows only 10 versions because PowerDNS implemented for *VERSION.BIND* requests from version 3.0 [26]. The latest version of PowerDNS Authoritative Server and PowerDNS Recursor are 3.3 and 3.5.2 in July 2013. Figure 5 reveals that the latest version of PowerDNS are mainly used.

4.3　Dnsmasq

Dnsmasq is a lightweight DNS forwarder and DHCP software for small network. Even though Dnsmasq isn't designed as a large scale resolver, many Dnsmasq servers are open resolvers shown in table 2.

Table2　Types of DNS Servers

Type of DNS	#	Total %	APNIC #	RIPE #	ARIN #	LACNIC #	AFRINIC #	other #
BIND 9.x	4268442	(14.1%)	806357	1530177	1126501	169268	121556	514583
†	1851362	(6.1%)	551458	781954	176399	94385	117906	129260
BIND 8.x	35218	(0.1%)	4588	21348	6663	974	32	1613
†	30444	(0.1%)	4202	18958	5186	854	31	1213
BIND 4.x	3486	(0.0%)	121	2751	440	43	0	131
†	2765	(0.0%)	93	2256	348	11	0	57
Dnsmasq	1308653	(4.3%)	692042	216273	75201	226880	32676	65581
†	1308381	(4.3%)	692026	216028	75196	226877	32676	65578
Nominum Vantio	968041	(3.2%)	553404	284852	20142	21205	70861	17577
†	967044	(3.2%)	552650	284782	20125	21200	70736	17551
Nominum ANS	687	(0.0%)	18	34	79	42	2	512
†	13	(0.0%)	2	0	0	11	0	0
PowerDNS	373588	(1.2%)	14215	329994	14360	2952	91	11976
†	372684	(1.2%)	14207	329116	14354	2952	91	11964
Unbound	71781	(0.2%)	16230	43507	6941	1510	1585	2008
†	23220	(0.0%)	3281	14398	4638	315	312	276
NSD	33933	(0.1%)	1731	11077	17182	322	13	3608
†	17	(0.0%)	5	5	2	1	0	4
Windows series	11698	(0.0%)	184	1077	85	10312	0	40
†	11342	(0.0%)	129	865	67	10257	0	24
can't detect	8281885	(27.3%)	4012525	2367711	429450	690618	279903	501678
†	7658656	(25.3%)	3911886	2118455	244682	670597	278183	434853
no version info	14927910	(49.3%)	3457029	4505928	1442348	4025325	699029	798251
†	12746062	(42.1%)	3050589	3465814	1179188	3919438	668399	462634
Total	30285322	(100.0%)	9558444	9314729	3139392	5149451	1205748	1917558
†	24971990	(82.5%)	8780528	7232631	1720185	4946898	1168334	1123414

†: open resolver
measurement on 5th and 6th of July, 2013

We totally found 86 software versions of Dnsmasq, and the figure 6 shows version distribution of it. The latest version of Dnsmasq is 2.66 in July 2013, but the version 2.66 doesn't appear in the figure.

4.4 Unbound and NSD

Unbound is caching and resolving name server software, NSD is authoritative name server software, and both of them have being developed by NLnet Labs [19]. Figure 7 and 8 show their version distribution respectively. We found 30 versions for Unbound and 42 versions for NSD in total. At the time of this measurement, the latest versions of Unbound and NSD were 1.4.20 and 3.2.15, respectively. We discovered that some beta version NSD servers are deployed on the Internet, such as 4.0.0b4 and 4.0.0_imp_5, from figure 8.

4.5 Nominum Vantio and ANS

Nominum Vantio and ANS are commercial caching and authoritative DNS servers developed by Nominum, respectively. These source codes are completely closed, so the latest versions of them aren't disclosed on the Internet. We infer that their latest versions are 5.3.3.1 and 5.3.1.0 from our measurement results at that time, respectively.

5 Domain Distribution of Open Resolver

We reversely looked up 30 million IP addresses to obtain FQDN. At first, we tried to reversely look up by using Unbound's library on the measurement computer, but we gave up this way because we estimated that it would take about 2 months to accomplish. Therefore, we then implemented the reverse lookupper, which asynchronously inquires FQDNs to Google public DNS, in figure 1 by using libevent. We accomplished this reverse look-up within

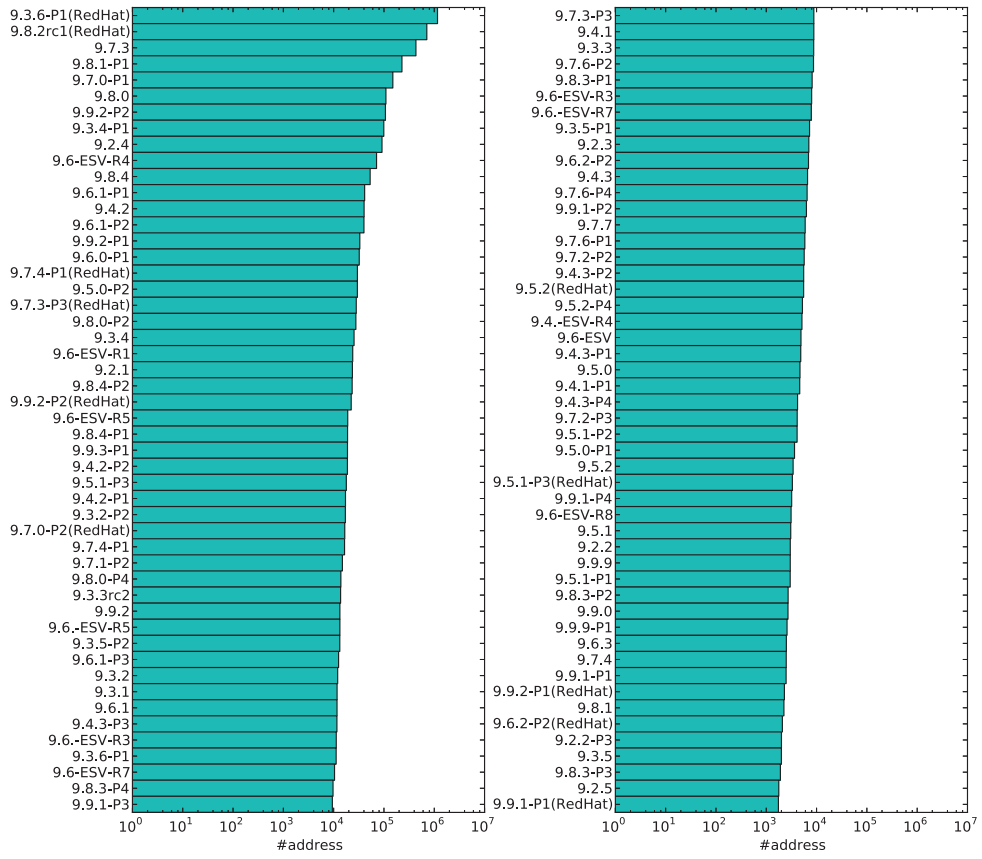

Figure2　Version Distribution of BIND 9.x Series (Top 100)

about 5 days with this implementation.

Figure 11 and 12 show 1st-to-3rd level domain distribution of all open resolvers and JP TLD's open resolvers, respectively.

We discovered spammer-favored domains, which are `163data.com.cn` and `hinet.net` reported by Craig A. Shue et al. [8], in table 11. To study more precisely, it should be compared with the population of domains and the population of spammer-favored domains, but this is outside the scope of this paper.

We found 381,387 addresses of JP TLD in total, and discovered that `ocn.ne.jp` is the worst holder of open resolvers in JP TLD. OCN managed by NTT Communications is the biggest and the most popular ISP in Japan [23]. The customer population should make OCN the worst holder.

6　Discussion

In this section, we discuss open resolvers and DNS the amplification attack, which is a DDoS attack abusing open resolvers as reflectors.

6.1　DNSSEC Considered Harmful

The DNS amplification attack can be launched because some types of DNS queries swell dozens of times when responding. For example, *ANY* query to `isc.org` and `ripe.net`, which are 64 and 65 bytes, are amplified to 3,245 and 2,669 bytes including IP and UDP header in August 2013, respectively. You can confirm this fact by running the following dig command: `$ dig any isc.org +bufsize=4096`.

Table 3 shows the details of DNS answer section of response for *ANY* query we obtained. It reveals that RRSIG, DNSKEY and NSEC records, which are records for DNSSEC [13, 14, 15], account for the majority of the response. In 2012, anonymous authors reported that some ISPs and governments, such as the Great Firewall of China, exploit AS level DNS injection attack for censorship [1]. DNSSEC can prevent the Internet users from such attack since it guarantees validity of DNS response. On the other hand, DNSSEC tremendously boosts the ef-

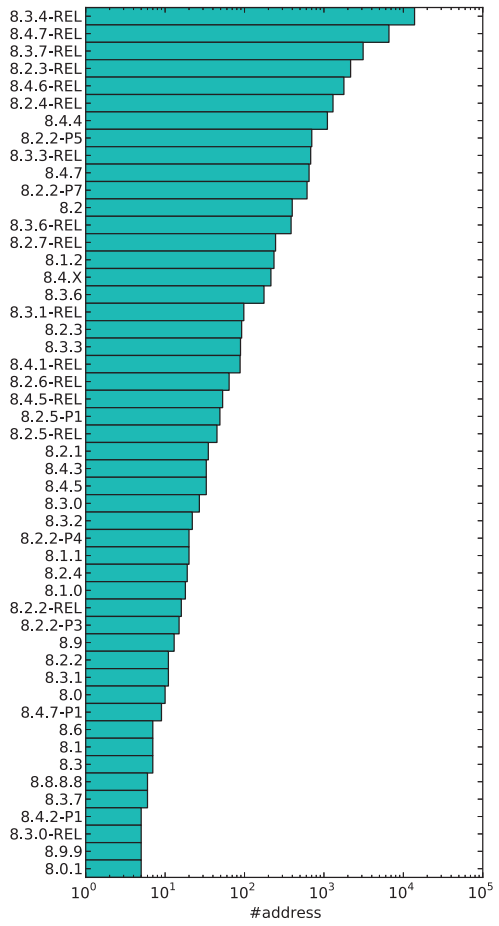

Figure3　Version Distribution of BIND 8.x Series (Top 50)

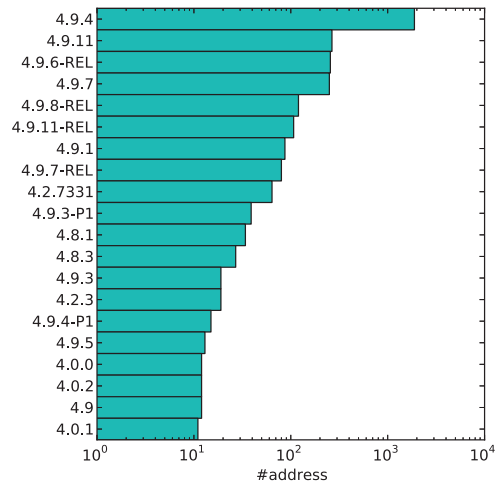

Figure4　Version Distribution of BIND 4.x Series (Top 20)

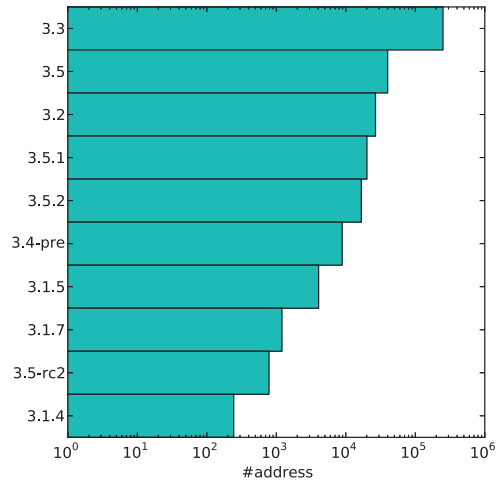

Figure5　Version Distribution of PowerDNS (Top 10)

ficiency of the DNS amplification attack.

6.2　Countermeasures

Updating DNS protocol is the most fundamental approach to cope with the open resolvers. If DNS protocol validates queriers, the DNS amplification attack can't be launched because it is performed by source address spoofing. Instead of UDP, applying TCP, which makes sure of sender when establishing connection by 3-way handshake, can prevent source address spoofing. However, TCP increases response time of DNS, even though it is required that DNS servers respond results as quickly as possible. For using TCP for DNS, fast TCP connection techniques would be helpful to reduce total latency of DNS query. TCP Fast Open [29] can reduce total round trip time by sending data to the peer before receiving ACK packet on 3-way handshake. ASAP [31], which adopts public key infrastructure to eliminate 3-way handshake, proposed by Wenxuan Zhou et al. can also reduce it.

Stopping every, about 25 millions, open resolvers on the Internet is another solution, but it is an unrealistic approach because the Internet is a distributed, autonomous and decentralized network. There is no centralized controller on the Internet, and even if some countries succeed in stopping open resolvers of them, open resolvers on spammer-favored domains will be still alive. Furthermore, as discussed in previous section, we discovered that DNS servers of obsoleted version still exist like BIND 4.x and 8.x. This fact implies that this solution isn't practical.

Applying egress filter by ISP is another way to disable source address spoofing attacks. It is also limited to solve the problem, but it is expected that appropriate egress filter mitigates the efficiency of this attack. If egress filter is

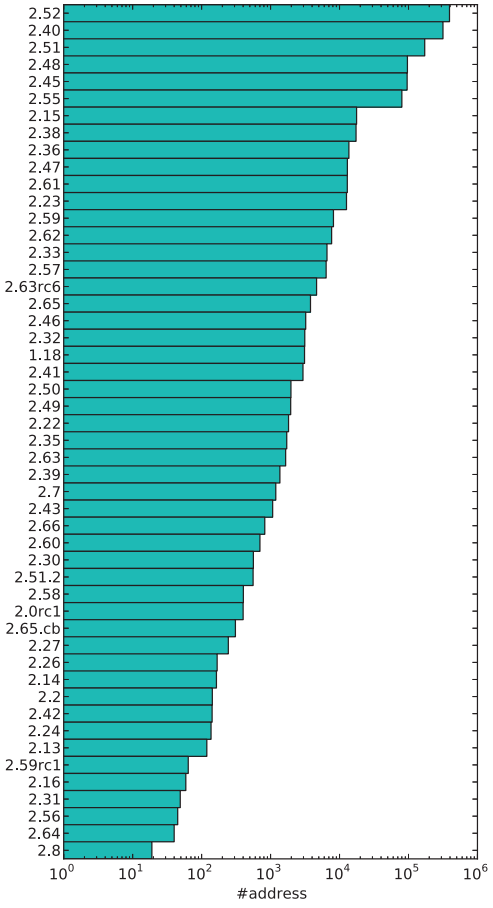

Figure6　Version Distribution of Dnsmasq (Top 50)

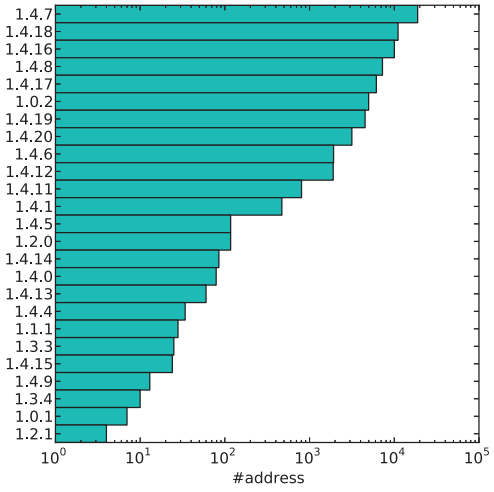

Figure7　Version Distribution of Unbound (Top 25)

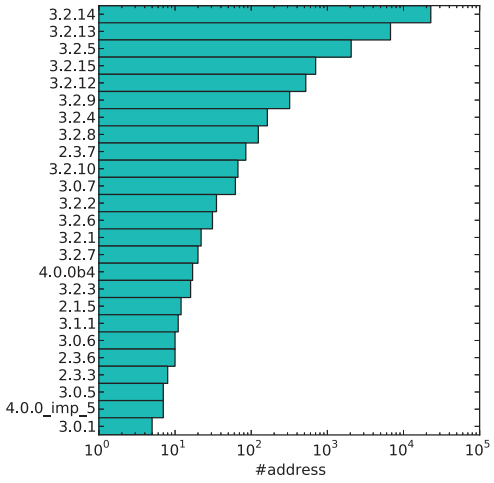

Figure8　Version Distribution of NSD (Top 25)

applied by many ISPs, source address spoofing attacks by botnets or script kiddies will be inefficient.

7　Conclusion and Future Work

This paper showed our measurement results of DNS servers, especially open resolvers, on the Internet. It revealed that there are about 30 million DNS servers, about 25 millions of which are open resolvers, and 7 millions of which respond software version request. We classified them by DNS server types and RIRs. The classification revealed that DNS servers of APNIC and RIPE NCC account for about 62.3 % of all DNS servers. It also revealed that obsoleted BIND 4.x and 8.x series are still alive, and RIPE NCC is the worst holder of them. In addition to this, we gave version distributions of each DNS type. It revealed that DNS server software versions have a wide distribution. The result implies that software once widely deployed isn't completely replaced to newer version. Furthermore, we reversely looked up 30 millions of address to obtain FQDN and gave 1st-to-3rd level domain distribution. It revealed there are many open resolvers on spammer-favored domains.

Furthermore, we discussed the DNS amplification attack, which abuses open resolvers as reflectors, and countermeasures of it. Based on that, we will study feasible approach toward such attacks in our future work.

Acknowledgment

We thank Prof. Yoichi Shinoda of JAIST. We are sure that we couldn't have accomplished our measurements without his help and support. We also thank Kunio Akashi of JAIST for supporting our measurement environment

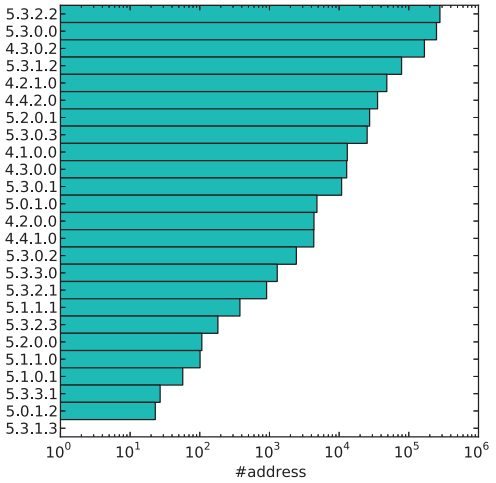

Figure9　Version Distribution of Nominum Vantio (All)

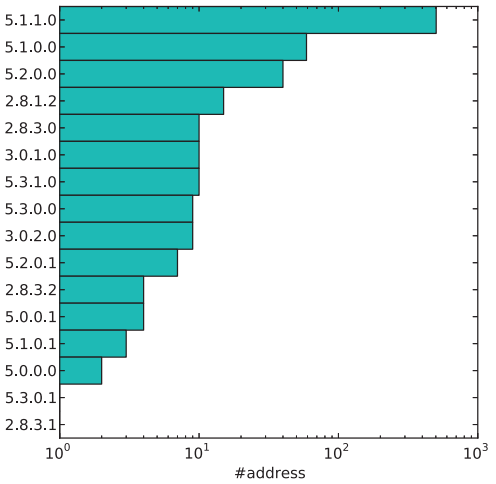

Figure10　Version Distribution of Nominum ANS (All)

setup.

References

[1] Anonymous. The Collateral Damage of Internet Censorship by DNS Injection. *SIGCOMM Computer Communication Review*, 42(3):21–27. http://conferences.sigcomm.org/sigcomm/2012/paper/ccr-paper266.pdf.

[2] Internet Systems Consortium — BIND. http://www.isc.org/downloads/bind/.

[3] BIND8 entering end of life. https://lists.isc.org/pipermail/bind-announce/2007-August/000222.html.

Table3　Details of DNS Answer Section of Response for *ANY* Query

	isc.org	ripe.net
RRSIG	1965	1304
DNSKEY	427	848
NSEC	53	38
SPF	112	-
TXT	181	-
NS	97	136
NAPTR	46	-
A	16	16
AAAA	28	28
MX	24	50
SOA	54	52
Total	3005	2472

(bytes)

[4] Internet Systems Consortium — BIND Software Status. http://www.isc.org/downloads/software-support-policy/bind-software-status/.

[5] Boost C++ Library. http://www.boost.org/.

[6] Catenaccio DPI. https://github.com/ytakano/catenaccio_dpi.

[7] CloudFlare - The DDoS That Knocked Spamhaus Offline (And How We Mitigated It). http://blog.cloudflare.com/the-ddos-that-knocked-spamhaus-offline-and-ho.

[8] Craig A. Shue and Minaxi Gupta and John J. Lubia and Chin Hua Kong and Asim Yuksel. Spamology: A Study of Spam Origins. In *Conference on Email and Anti Spam (CEAS)*, 2009.

[9] DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION (RFC 1035). http://tools.ietf.org/rfc/rfc1035.txt.

[10] US-CERT Alert(TA13-088A) DNS Amplification Attacks. http://www.us-cert.gov/ncas/alerts/TA13-088A.

[11] Dnsmasq - a DNS forwarder for NAT firewalls. http://www.thekelleys.org.uk/dnsmasq/doc.html.

[12] DOMAIN NAMES - CONCEPTS and FACILITIES (RFC 1035). http://tools.ietf.org/rfc/rfc882.txt.

[13] DNS Security Introduction and Requirements (RFC 4033). http://tools.ietf.org/rfc/rfc4033.txt.

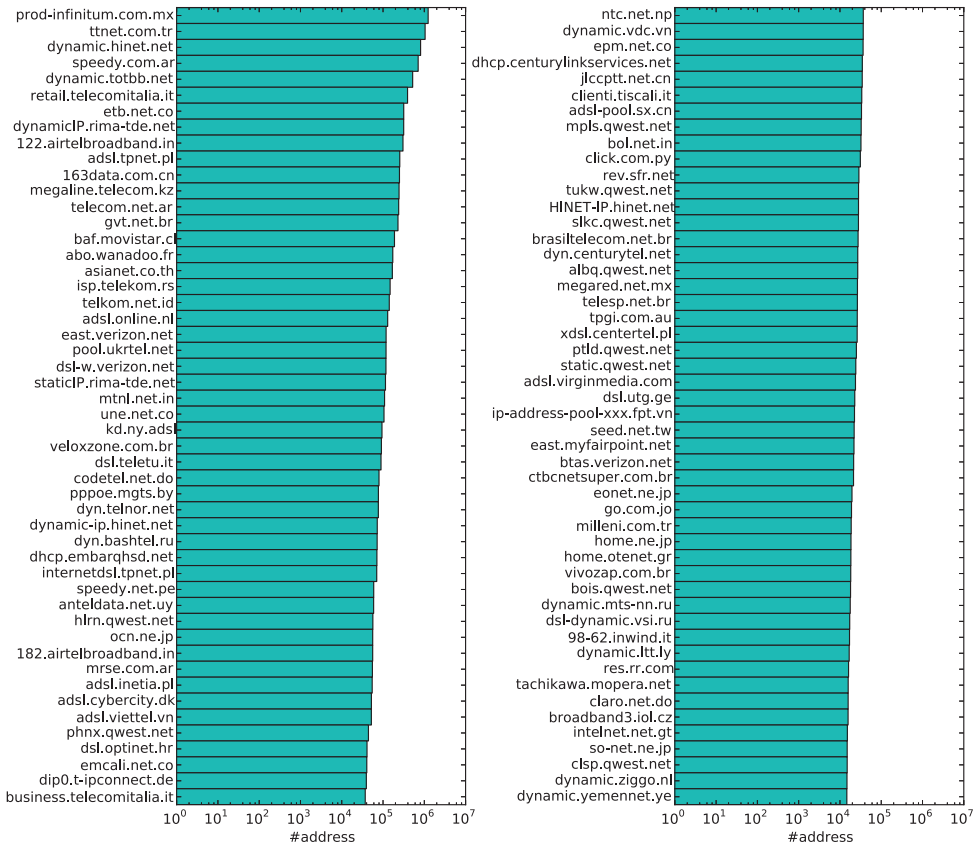

Figure11　1st-to-3rd Level Domain Distribution of Open Resolver (Top 100)

[14] Resource Records for the DNS Security Extensions (RFC 4034). http://tools.ietf.org/rfc/rfc4034.txt.

[15] Protocol Modifications for the DNS Security Extensions (RFC 4035). http://tools.ietf.org/rfc/rfc4035.txt.

[16] Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. In *OSDI*, pages 137–150. USENIX Association, 2004.

[17] libevent. http://libevent.org/.

[18] MongoDB. http://www.mongodb.org/.

[19] nlnetlabs.nl :: Home ::. http://nlnetlabs.nl/.

[20] Authoritative DNS — Nominum. http://www.nominum.com/products/core-engines/authoritative-dns/.

[21] Vantio Caching DNS — Nominum. http://www.nominum.com/products/core-engines/caching-dns/.

[22] nlnetlabs.nl :: Name Server Daemon (NSD) ::. http://www.nlnetlabs.nl/projects/nsd/.

[23] OCN Top Page. http://www.ocn.ne.jp/.

[24] Open Resolver Project. http://openresolverproject.org/.

[25] Welcome to PowerDNS. https://www.powerdns.com/.

[26] PowerDNS Authoritative Server 3.0 Release Notes. http://doc.powerdns.com/html/changelog.html#changelog-auth-3-0.

[27] IANA IPv4 Address Space Registry. http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.txt.

[28] Steve Santorelli. The global open resolver picture. http://www.securityacts.com/securityacts03.pdf#page=29.

[29] TCP Fast Open (IETF Draft). https://tools.ietf.org/html/draft-ietf-tcpm-fastopen-04.

[30] Unbound. http://unbound.net/.

[31] Wenxuan Zhou, Qingxi Li, Matthew Caesar, and Brighten Godfrey. ASAP: a low-latency transport

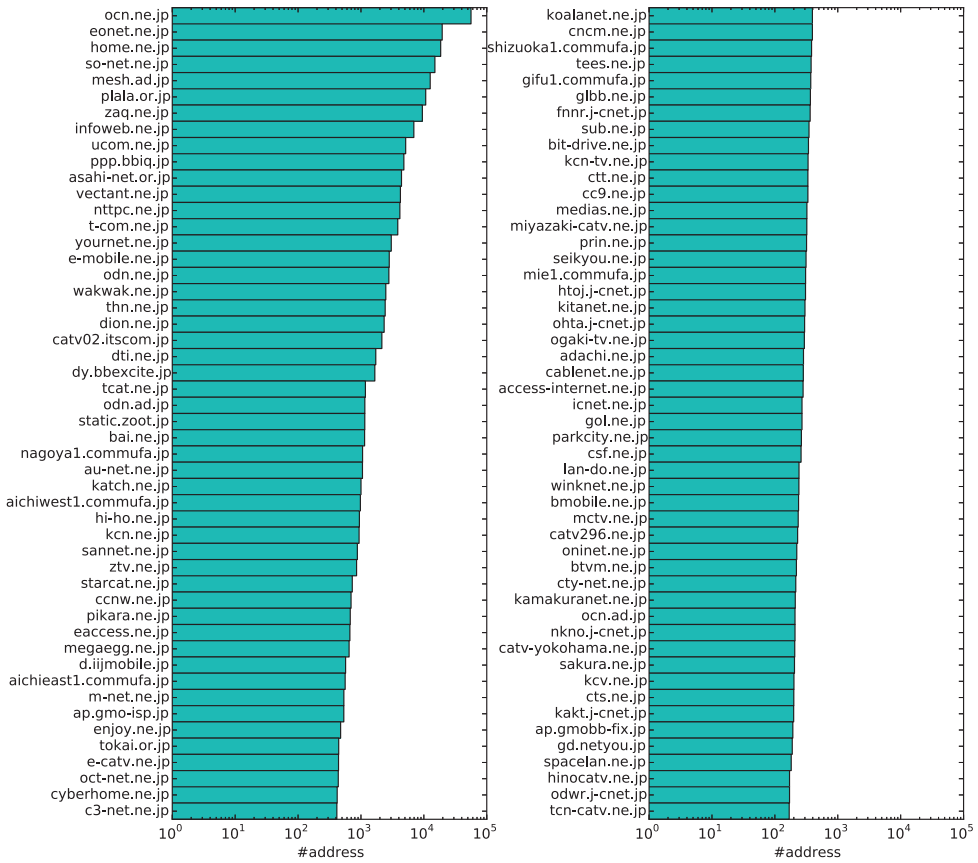

Figure12　1st-to-3rd Level Domain Distribution of Open Resolver in JP TLD (Top 100)

layer. In Kenjiro Cho and Mark Crovella, editors, *CoNEXT*, page 20. ACM, 2011.