A survey on the status of measures against IP fragmentation attacks on DNS

Kenya Ota[1,a)] Tsunehiko Suzuki[1,b)]

Abstract: The risk of DNS cache poisoning attacks using IP fragmentation was presented by Herzberg and Shulman in 2012 and 2013. And we showed that the attacks are feasible, and several open-source implementations were still affected by the attacks. In the wake of our proposal, measures to major open-source implementations for ignoring NS records in Authority or Additional sections of negative response at DNS cache server, and for ignoring Path MTU Discovery at DNS authoritative server were taken. Also, DNS flag day 2020 is planned to take measures against fragmentation attacks such as reducing default EDNS buffer size. If the authoritative servers that manage TLDs or multiple zones have not been taken measures, this attack increases the risk of massive hijacking at once. In this research, we survey whether authoritative servers that manage TLDs can be affected by the attacks.

Keywords: DNS, cache poisoning, IP fragmentation, security

1. Introduction

Domain Name System (DNS) has some threats, such as cache poisoning attacks. Cache poisoning is an attack that the attacker sends spoofed DNS response to inject fake resource records (RRs) to the full-service resolver's cache. As a measure against these attacks, DNS Security Extensions (DNSSEC) has been standardized. DNSSEC enables verification of origin and integrity of the response by validating digital signatures based on public-key cryptography.

DNSSEC-signed TLDs exceed 90% as of August 6, 2019 [1]. However, most zones that registered with TLDs have not signed. Moreover, according to APNIC's statistics, DNSSEC full-validating resolvers are around 20% [*1].

In such a situation, Herzberg and Shulman presented a new cache poisoning attack concept using IP fragmentation (fragmentation attacks) in 2012 [2] and 2013 [3]. And Hlaváček presented that the attacker can trigger IP fragmentation using Path MTU Discovery (PMTUD) [4]. Based on those, We reproduced the concept and confirmed that fragmentation attacks are feasible, and several open-source implementations were still affected by the attacks [5].

In the wake of our proposal, some major open-source resolver implementations took measures to ignore NS RRs in the Authority or Additional sections of negative response [6], [7]. Also, authoritative server implementations took measures to ignore PMTUD [8], [9], [10]. DNS flag day 2020, which focuses on the problems caused by IP fragmentation, is planned [11], and the EDNS buffer size value recommended by the DNS community was discussed [12].

The attacker can abuse an un-measured authoritative server for exploiting fragmentation attacks. Also, if the ISP or organization uses an un-measured full-service resolver, that resolver may be poisoned by the attacks. Particularly, if the TLDs or DNS hosting service operators use un-measured authoritative servers, the attacks increase the risk of massive hijacking at once. Therefore, in this research, we surveyed the authoritative servers that manage TLDs to determine whether they can be affected by the attacks. We report the result that more than half of the TLDs were affected, and it became clear that measures were not progressing from August to October.

2. Related Work

Research on IP fragmentation using PMTUD and survey of the fragmentation status of DNS responses depending on setting the EDNS buffer size are being conducted.

Göhring et al. investigated how common PMTUD is in actual communications using a data set from the Center for Applied Internet Data Analysis (CAIDA) [13]. That indicated approximately 95.7% of the Next-Hop MTU value in PMTUD is in the range of 1350 to 1500 bytes. Also, they investigated whether or not changing PMTU using PMTUD is possible for a total of 5000 domains of Alexa Top 1M's top 4000 domains and 1000 domains from 100,000. As a result, it was found that about 80% of the servers were reduced by less than 600 bytes.

Fujiwara surveyed the fragmentation status of the response of Alexa top 1M domain [14]. That survey queried domain name for A and AAAA RRs and compared the DNS

1 Graduate School of Engineering, Chukyo University
a) t31903m@m.chukyo-u.ac.jp
b) tss@suzuki.sist.chukyo-u.ac.jp

*1 https://stats.labs.apnic.net/dnssec/

fragmentation status when the EDNS buffer size was set to 4096 bytes and 1220 bytes. That showed when the EDNS buffer size is set to 4096 bytes, 64334 packets (about 0.3% of the total packets, 2438 IPv4 addresses) were fragmented. In contrast, in the case of 1220 bytes, the number was reduced to 26 packets.

Brandt et al. showed a technique for issuing certificates illegally from Certificate Authorities (CA) that issues Domain Validation (DV) certificates [15]. Cache poisoning is performed on the CA resolver, and the attacker issues a certificate by illegally proceeding with the e-mail authentication procedure. In addition to proposing DV improvement methods in this paper, it is shown that IP fragments are prevented, and DNSSEC is fully supported as countermeasures against this attack. Also, Let's Encrypt changed the EDNS buffer size to 512 bytes based on this research [16].

3. Fragmentation Attacks

In this section, we explain the concept, attack examples, and measures of fragmentation attacks.

3.1 Concept

Fragmentation attacks abuse the IP fragmentation reassembly process. On DNS and UDP, fragmented IP packets excluding the first fragment do not contain UDP header (source port number) and DNS Header section (transaction ID, query name, and count of RRs in each section). Hence, the attacker attempts to tamper with a legitimate DNS response by replacing the second or following fragments. The impact of the attacks depends on the full-service resolver implementations and cached data.

We show the attack procedure targeted to the open resolver in **Fig. 1**. The attacker considers the DNS query that the response causes IP fragmentation, and spoofed RRs will be cached. Next, the attacker sends some spoofed second fragment packets to the victim resolver. The source IP address of these packets must be set to the target authoritative server, and IP Identification (IP-ID) is set as random. After sending fragment packets, the attacker sends a DNS query to the victim resolver to trigger name resolution. If there is a spoofed second fragment that matches the IP-ID of the legitimate first fragment, the OS will reassemble these packets. Then, if the resolver accepts the reassembled response, the attack will succeed.

When the fragmented packets are received, the OS is buffering these packets until the host receives all fragments. On the Linux kernel, the buffer size is 64 packets by default, and the timeout is 30 seconds. The reassembly process does not depend on packet arrival orders. Hence, the attacker can send spoofed fragments up to the limit of buffer before sending a DNS query.

To performing the reassembly process, several fields such as IP-ID, UDP checksum must match with the legitimate DNS response. On DNS, if the zone information and server configuration are consistent, it can be expected that the authoritative server returns the same response to the same

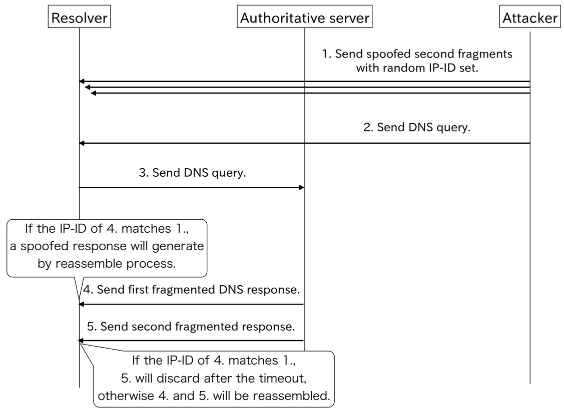

Fig. 1 Procedure of IP fragmentation attacks.

query.

In addition, UDP checksum is 2 bytes value, which is calculated as one's complement sum of UDP pseudo header and payload and one's complement. UDP checksum can be adjusted by changing the TTL value in the RRs or by using the EDNS padding option. Therefore, the attacker can obtain a legitimate response and adjust UDP checksum. The attacker only needs to predict IP-ID value and the entropy of DNS response decreases to 16 bits on IPv4.

3.2 Path MTU Spoofing

Most DNS response sizes are less than 1500 bytes. Hence the response is less likely causing IP fragmentation. Exploiting attacks need to trigger fragmentation. It is also beneficial for an attacker to adjust the fragmentation position to the boundaries of the sections or RRs in the DNS message. These can be executed by PMTU spoofing that exploits PMTUD.

PMTUD is a mechanism to suppress fragmentation on the path by searching for an MTU on the path and cause fragmentation by the sender. That is standardized in RFC 1191 on IPv4, and RFC 1981, 8210 on IPv6. On IPv6, IP fragmentation must be caused by the end node.

We show an example of the PMTU spoofing attack on IPv4 in **Fig. 2**. The attacker sends an ICMP echo request to the authoritative server which force to cause fragmentation. The source IP address is set to the target resolver. Then, the attacker sends ICMP type=3 (destination unreachable), code=4 (fragmentation needed and DF set) message to the authoritative server. This packet may be called Packet Too Big (PTB). The Next-Hop MTU value in PTB can be set to an arbitrary size that the attacker wants to cause fragmentation. If the authoritative server accepts the PTB packet, the server will cause fragmentation for packets destined for the resolver.

As a result of our confirmation on IPv4, Arch Linux (Linux Kernel 5.1.9) accepts PTB, and the PMTU can be decreased to 552 bytes. In contrast, FreeBSD 12.0 ignores the PTB packet. Therefore, if the authoritative server running on Linux, the server may be abused for the attacks.

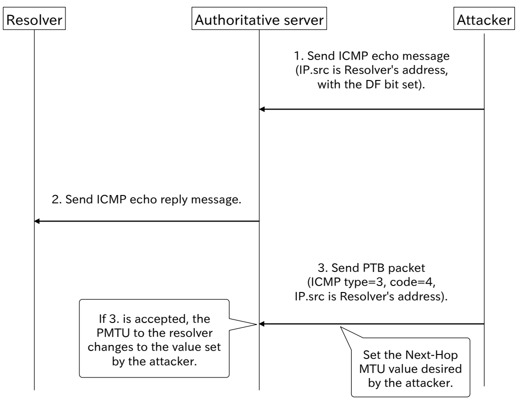

Fig. 2 Procedure of PMTU spoofing.

Fig. 3 Example of legitimate negative response.

Fig. 4 Example of spoofed negative response.

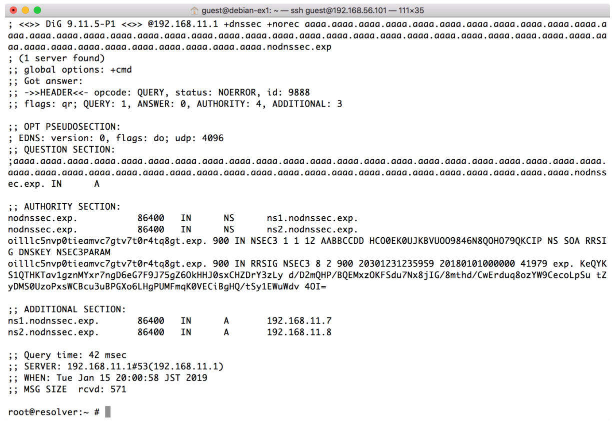

Fig. 5 Example of legitimate delegation response.

3.3 Attack Vectors

3.3.1 Negative Response Replacement Attack

DNSSEC uses NSEC or NSEC3 RRs to prove authenticated denial of existence. Consider the case where a client queries a non-existence name with the DO bit set. If the zone that has received the query is DNSSEC-signed, the authoritative server returns negative response, including NSEC/NSEC3 and RRSIG RRs in the Authority section. The negative response replacement attack replaces these RRs with fake NS RRs.

We show a legitimate response example in **Fig. 3** and a spoofed response example in **Fig. 4**. In this attack example, we replace the last RRSIG RR in the Authority section with NS RR and adjusted UDP checksum and message length by the EDNS padding option.

This attack is based on ranking data described in RFC 2181 and the negative response format that showed in RFC 2308. Ranking data defines the trustworthiness of data. A negative response is an authoritative answer, and NS RRs in the Authority section of that response is usually more trustworthy than the cached record. In addition, RFC 2308 describes some response examples that include NS RRs in the Authority section. If the resolver implementation that is compliant with these RFCs and accepts NS RRs in the Authority section, this attack will succeed.

If the full-service resolver is queried for a name or record that does not exist in the cache, the resolver will query the authoritative server. For this reason, an attacker can cause a query by selecting a random label that does not exist in the cache and can attack continuously. However, it is necessary to select a label or fragmentation position so that the payload of the second fragment is constant.

With this attack, domain hijacking is possible with the non-validating resolvers. Also, subdomain injection may be possible even if the resolver is validating when the target zone uses NSEC3 Opt-Out [2], [3].

3.3.2 Delegation Response Replacement Attack

When the child zone is delegated to other organizations, the parent zone returns a delegation response. Delegation response includes NS RRs in the Authority section. Also, if there are A/AAAA RRs corresponding to NS RRs and the name is the in-bailiwick name of the parent zone, the response includes these RRs as glue records. The delegation response replacement attack replaces these glue records with fake A/AAAA RRs. This attack is replacing these glue records to induce fake authoritative servers.

We show a legitimate response example in **Fig. 5** and a spoofed response example in **Fig. 6**. In these examples, we show a delegation response from the DNSSEC-signed zone to an unsigned zone.

This attack is based on DNSSEC specifications that non-authoritative RRs do not have RRSIG RRs. Since the glue records are non-authoritative information, there are

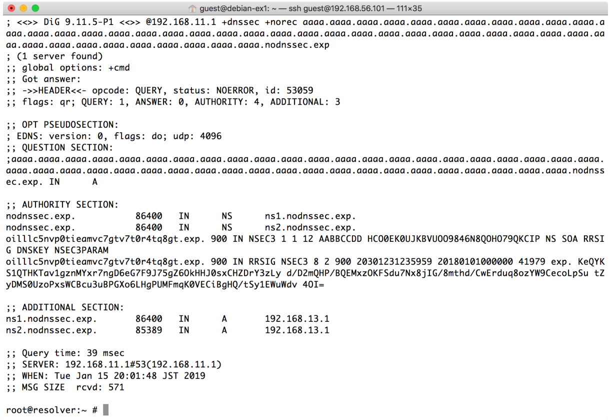

```
guest@debian-ex1: ~ — ssh guest@192.168.56.101 — 111×35
; <<>> DiG 9.11.5-P1 <<>> @192.168.11.1 +dnssec +norec aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.a
aaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aa
aa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.nodnssec.exp
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 53059
;; flags: qr; QUERY: 1, ANSWER: 0, AUTHORITY: 4, ADDITIONAL: 3

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags: do; udp: 4096
;; QUESTION SECTION:
;aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.
aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.aaaa.nodnss
ec.exp. IN      A

;; AUTHORITY SECTION:
nodnssec.exp.           86400   IN      NS      ns1.nodnssec.exp.
nodnssec.exp.           86400   IN      NS      ns2.nodnssec.exp.
oilllc5nvp0tieamvc7gtv7t0r4tq8gt.exp. 900 IN NSEC3 1 1 12 AABBCCDD HCO0EK0UJKBVUO09846N8QOHO79QKCIP NS SOA RRSI
G DNSKEY NSEC3PARAM
oilllc5nvp0tieamvc7gtv7t0r4tq8gt.exp. 900 IN RRSIG NSEC3 8 2 900 20301231235959 20180101000000 41979 exp. KeQYK
S1QTHKTav1gznMYxr7ngD6eG7F9J75gZ6OkHHJ0sxCHZDrY3zLy d/D2mQHP/BQEMxzOKFSdu7Nx8jIG/8mthd/CwErduq8ozYW9CecoLpSu tZ
yDMS0UzoPxsWCBcu3uBPGXo6LHgPUMFmqK0VECiBgHQ/tSy1EWuWdv 4OI=

;; ADDITIONAL SECTION:
ns1.nodnssec.exp.       86400   IN      A       192.168.13.1
ns2.nodnssec.exp.       85389   IN      A       192.168.13.1

;; Query time: 39 msec
;; SERVER: 192.168.11.1#53(192.168.11.1)
;; WHEN: Tue Jan 15 20:01:48 JST 2019
;; MSG SIZE  rcvd: 571

root@resolver:~ #
```

Fig. 6 Example of spoofed delegation response.

no RRSIG RRs corresponding to glue records, even if the zone is DNSSEC-signed. Therefore, even if the parent zone is DNSSEC-signed and the resolver is validating DNSSEC response, the resolver may cache spoofed glue records. Note that when the delegated zone is DNSSEC-signed, the DNSSEC-validating resolver will handle the response as SERVFAIL.

In Fig. 5 and Fig. 6, NSEC3 and RRSIG RRs are used to increase the response size. However, when a more significant number of NS RRs in the Authority section or long label is used to NS RRs, these increase the possibility of fragmentation.

3.3.3 Other Attack Vectors

It is conceivable to use a response from a zone in which a wildcard is set (e.g., the owner of the RR is `*.example.jp.`). A wildcard can always cause a query by randomly querying a domain name that does not exist in the cache, as in the case of a negative response.

Since the response message size, including a large number of NS RRs and long TXT RRs, exceeds the MTU and causes fragmentation, it can be used for attacks.

Another possible attack is to contaminate the sibling domain's NS and A/AAAA RRs (e.g., delegation response when querying `small-is-beautiful.jp.` to `jp.`). In this case, even if the response of the own domain is small, the response becomes weak if the response of the sibling domain is massive. The attacker could prepare the zone.

3.4 Measures

In the wake of our proposal, Unbound 1.8.2 and later [6], and PowerDNS Recursor 4.2.0 and later [7] have taken measures to ignore NS RRs in Authority section of negative response. In our confirmation, BIND and Knot Resolver do not affect the negative response replacement attack. Also, the Linux Kernel 3.15 and later have `IP_PMTUDISC_OMIT` socket option that ignores PMTUD. NSD 4.1.27 and later [8], Knot DNS 2.8.2 and later [9], and PowerDNS Authoritative Server 4.2.0-rc2 and later [10], this option is used to avoid fragmentation. BIND has been used for this option before our proposal [17]. Note that, on IPv4, FreeBSD avoids PTB, so it seems any implementations are unaffected.

Effective countermeasures include avoiding fragmentation by reducing the EDNS buffer size and dropping fragmented DNS responses using a firewall. These measures are also mentioned in [14], [18], [19]. The recommended value of buffer size was discussed [12], and the value such as 1220, 1232, 1280 bytes are listed as candidates in the DNS community. At the time of writing this paper (November 7, 2019), the recommended value in [11] is 1232 bytes. However, since even smaller responses can be attacked, we propose 512 bytes the same as usual.

Other measures include reducing query name length with QNAME Minimisation, caching each zone's NS RRs as the authoritative data, and deploying DNSSEC completely. However, Herzberg and Shulman said, "incremental DNSSEC deployment is vulnerable to our cache poisoning attacks" in [2].

Also, if the full-service resolver does not perform DNSSEC validation, the resolver sends queries without the DO bit to reduce the response size. However, in RFC 4035, the DNSSEC-compatible resolver must always turn on the DO bit, and some implementations such as Unbound cannot turn off the DO bit with configuration [*2].

4. Survey

4.1 Objectives

As shown in section 3, it is possible to easily tamper DNS messages by causing IP fragmentation using PMTU spoofing and the response with the DO bit set. The impact of fragmentation attacks becomes more serious when the shared authoritative server that manages TLD or multiple zones is targeted. Particularly, TLDs are DNSSSEC-signed more than 90% [1], and the response size is more significant than the unsigned zone. Therefore, in the un-measured full-service resolver, the TLD hijack is possible if the resolver does not perform DNSSEC verification using the negative response replacement attack. Even if the resolver validates the DNSSEC signatures, subdomain injection may be successful [2], [3].

From the above, it is said that if the TLD's authoritative server that affects PMTU spoofing is more dangerous. Thus the objective of this survey is to clarify the countermeasure situation in TLDs.

4.2 Method

We scanned 1387 TLDs that DS RR is registered in the root zone as of August 6, 2019 [1]. We conducted the scan in August and October 2019. In the scan, we sent PTB packets and DNS queries to each TLD's authoritative servers. Then, we inspected whether the authoritative servers accepted PTB and replied fragmented DNS responses based on that MTU value. The scan was based on IP addresses, the authoritative server that manages multiple TLDs was scanned any one of the TLDs. We use FreeBSD 12.0 and scapy 2.4.0 to send and analyze packets, and tcpdump to

*2 Unbound can be turned off the DO bit by changing EDNS_DO in util/net_help.h.

capture packets.

4.2.1 PMTUD Scan

We scanned whether the PTB packet can change the PMTU value. The procedure is as follows.

(1) Send a 1454-byte ICMP echo message to the TLD's authoritative server with the DF bit set.
(2) Check the ICMP echo reply message from the authoritative server.
(3) If there is a response from the authoritative server, send PTB with Next-Hop MTU set to 68 bytes for ICMP echo reply, then send ICMP echo message again.
(4) check the ICMP echo reply again.

The above was repeated five times to check whether the ICMP echo reply was fragmented. When the authoritative server returned a fragmented response, the packet size of the first fragment was also recorded. The timeout was set to 2 seconds, and if there was no response during that time, it was judged as "noreply". Note that when the authoritative server that did not reply ICMP echo request, we do not execute step 3 and 4.

4.2.2 DNS Fragmentation Scan

We scanned whether the DNS response can be fragmented by the PMTU value set by the PTB packet. The procedure is as follows.

(1) query a non-existence name with the DO bit set and the EDNS buffer size set to 2048 bytes after the PMTUD scan.
(2) check the DNS response.

The above was repeated five times to check whether the DNS response was fragmented. A DNS response was recorded as fragmented when the first fragment was less than or equal to the size recorded by the PMTUD scan. When the authoritative server returned a fragmented response, the packet size of the first fragment was also recorded. The timeout was set randomly between 2 and 5 seconds, and if there was no response during that time, it was judged as "noreply". This scan also recorded EDNS buffer size values for all responses. Note that the NS RRs that could not be resolved are not scanned.

4.3 Result

We show the number of NS RRs each TLD has in **Fig. 7**. In the scan conducted in August, the average number of NS RRs registered in each TLD was approximately 4.80, including those that could not be resolved. Of these, a total of 3151 hosts were used as authoritative servers for name resolution in all TLDs. In October, the average number of NS RRs was approximately 4.78 RRs, and a total of 3127 hosts were used as authoritative servers for name resolution in all TLDs. The number of NS RRs that could not be resolved was 45 in August and 42 in October. As of October, 3 TLDs listed in [1] were removed from the root zone. Also, in both August and October, 3107 hosts were used, 44 were not used, and 20 hosts were added.

We show the scanned result for each host in **Table 1**. As a result of the scan conducted in August, 3071 hosts

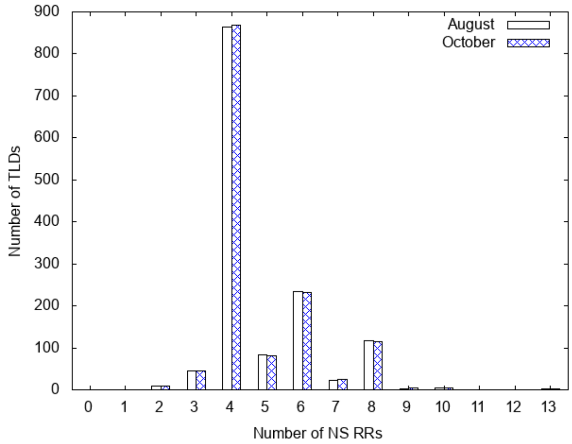

Fig. 7 Number of NS RRs registered for each TLD.

Table 1 Fragmentation status of ICMP and DNS responses per host.

	DNS frag					
	August			October		
ICMP frag	Yes	No	noreply	Yes	No	noreply
Yes	1792	328	3	1759	334	3
No	52	896	0	52	902	1
noreply	0	75	5	0	71	5

Table 2 Correspondence in DNS fragmentation status for each host in August and October.

		October		
		Yes	No	noreply
August	Yes	1635	184	2
	No	164	1114	0
	noreply	1	0	7

$(1792+328+3+52+896+0)$ replied ICMP echo request, and the packets of 2123 hosts $(1792 + 328 + 3$, approximately 67.4%) are fragmented. On the other hand, 3051 hosts $(1759+334+3+52+902+1)$ replied ICMP echo request, and 2096 hosts $(1759+334+3$, approximately 67.0%) sent fragmented reply in October scan. The packet length of all first ICMP fragment was 548 bytes both August and October. As a result of DNS scan, we got responses from 3143 hosts $(1792+52+0+328+896+75)$, and 1844 hosts $(1792+52+0$, approximately 58.5%) replied fragmented responses in August. Whereas in October, we got responses from 3118 hosts $(1759+52+0+334+902+71)$, and 1811 hosts $(1759 + 52 + 0$, approximately 57.9%) replied fragmented responses.

We show the correspondence in DNS fragmentation status for each host in August and October in **Table 2**. 2756 hosts $(1635+1114+7$, approximately 88.7%) did not change the status.

We show the results of summarizing the number of hosts that returned fragmented DNS responses by PTB for each TLD in **Fig. 8**. The number of fragmented hosts increased from August to October. **Fig. 9** shows the percentage of hosts that returned the fragmented DNS responses of each TLD in August, and **Fig. 10** shows the results in October. **Fig. 11** shows that the change in fragmented NSs each TLD between August and October. The number of hosts increased by 523 TLDs (approximately 37.8% of all TLDs),

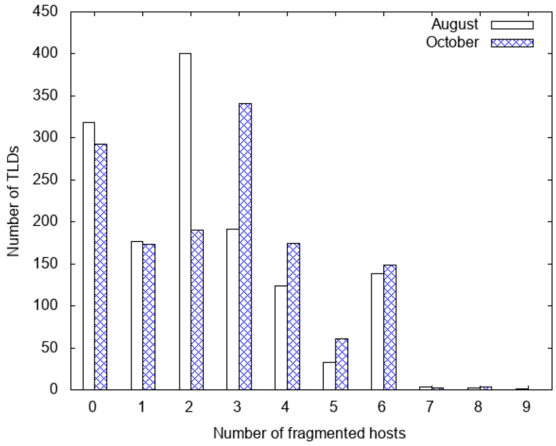

Fig. 8 Number of hosts that return fragmented responses due to the influence of the PTB of each TLD.

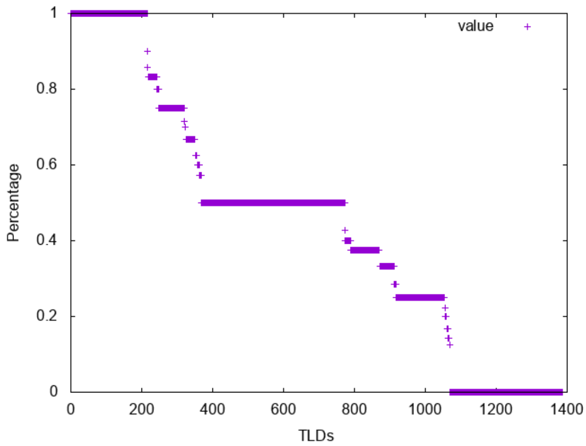

Fig. 9 Percentage of hosts returning fragmented DNS responses in NS RRs for each TLD in August.

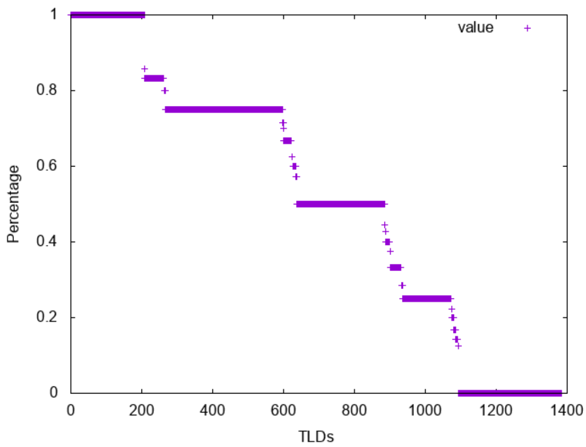

Fig. 10 Percentage of hosts returning fragmented DNS responses in NS RRs for each TLD in October.

of which one host increased by 453 TLDs (approximately 32.7%). It was 134 TLDs that decreased (approximately 10.0%), and 727 TLDs (approximately 52.5%) that did not change.

We show the results of totaling the EDNS buffer size of all responses in **Table 3**. Almost all responses were set to 4096 bytes. In addition, the number of responses with 1232 bytes set increased by 35.

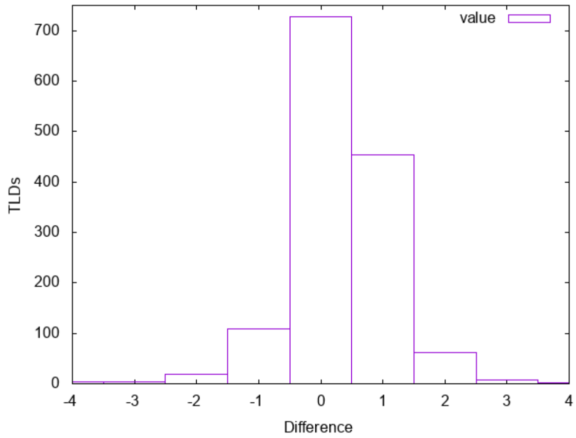

Fig. 11 Change in the number of hosts that return fragmented DNS responses from August to October.

Table 3 Number of responses for each EDNS buffer size value.

	count			
	August		October	
bufsize	all	frag	all	frag
512	35	0	35	0
1220	15	0	15	0
1232	30	0	65	0
1280	20	5	20	5
1432	18	10	15	10
1450	5	1	5	1
1472	3	0	2	0
1480	5	0	5	0
1680	5	5	5	5
4096	15540	5663	15393	5619
32768	5	0	5	0

4.4 Discussion

Table 1 shows that more than half of the hosts return fragmented DNS responses. Moreover, Table 1 and Table 2 shows that it can be said that there is almost no change in the overall trend. Fig. 11 shows that an increase in the number of hosts that return fragmented responses with approximately 38.7% TLDs. Most of those increased by one host. In order to distribute the load on route servers and TLDs with many accesses, anycast is used in which multiple hosts respond to the same IP address depending on the communication path. It can be considered that these results are affected by load balancing because there are few changes when viewed from each host. Moreover, there are some hosts where only the DNS response is fragmented without the ICMP fragmentation.

Since the number of TLDs that return fragmented DNS responses has hardly decreased, it can be said that the measures at the authoritative server have not progressed. Furthermore, it is a critical situation. In addition, there is a possibility that fragmentation attacks can be performed on more than half of the TLDs, which is a critical situation. Therefore, on the full-service resolver, it must take measures, as shown in section 3.4. The authoritative server should also take measures immediately.

From the results in Table 3, the EDNS buffer size has hardly changed from August to October. And it is thought that there are many hosts that use the default values (4096 bytes) of each implementation as they are. Other than 4096 bytes, most hosts set 512 bytes in August, but 1232 bytes were the largest in October. That may have been influenced by many opinions recommending 1232 bytes during discussions in the DNS community [11], [12], but the relationship is not clear.

5. Conclusion

We investigated the fragmentation status of ICMP and DNS responses by PTB for TLDs in August and October. More than half of the hosts returned DNS responses that were fragmented by PTB, and TLDs that used more than half of the affected authoritative servers showed a slight decrease. It has become clear that measures against fragmentation attacks on the authoritative server that manages TLDs have not progressed. Furthermore, it is a critical situation that requires measures in the full-service resolver side. TLDs are expected to take immediate measures.

References

[1] ICANN Research: *TLD DNSSEC Report (2019-08-06 00:02:39)*, `http://stats.research.icann.org/dns/tld_report/archive/20190806.000101.html`, (2019).
[2] Amir Herzberg and Haya Shulman: *Fragmentation Considered Poisonous*, CoRR (2012).
[3] Amir Herzberg and Haya Shulman: *Fragmentation Considered Poisonous, or: one-domain-to-rule-them-all.org*, In Proceedings of IEEE Conference on Communications and Network Security, pp. 224-232, (2013).
[4] Tomáš Hlaváček: *IP fragmentation attack on DNS*, `https://ripe67.ripe.net/presentations/240-ipfragattack.pdf`, RIPE 67, (2013).
[5] Kenya Ota and Tsunehiko Suzuki: *A proof of concept and countermeasures for 1st-fragment piggybacking attacks*, The 81st National Convention of IPSJ, pp. 443-444, (Japanese), (2019).
[6] NLnet Labs: *Unbound 1.8.2 released*, `https://www.nlnetlabs.nl/news/2018/Dec/04/unbound-1.8.2-released/`, (2018).
[7] PowerDNS: *[WIP] rec: Skip NS records in authority/additional sections of NXD / NODATA #7258*, `https://github.com/PowerDNS/pdns/pull/7258`, (2018).
[8] NLnet Labs: *[nsd-users] NSD 4.1.27 released*, `https://open.nlnetlabs.nl/pipermail/nsd-users/2019-March/002666.html`, (2019).
[9] CZ.NIC: *Version 2.8.2 – Knot DNS*, `https://www.knot-dns.cz/2019-06-05-version-282.html`, (2019).
[10] PowerDNS: *Changelogs for 4.2.x*, `https://doc.powerdns.com/authoritative/changelog/4.2.html#change-4.2.0-rc2`, (2018).
[11] dns-violations: *DNS flag day 2020*, `https://dnsflagday.net/2020/index.html`, (2019).
[12] dns-violations: *flag day 2020: Recommended EDNS buffer size #125*, `https://github.com/dns-violations/dnsflagday/issues/125`, (2019).
[13] Matthias Göhring, Haya Shulman, and Michael Waidner: *Path MTU Discovery Considered Harmful*, In Proceedings of IEEE 38th International Conference on Distributed Computing Systems, pp. 866-874, (2018).
[14] Kazunori Fujiwara: *Measures against DNS cache poisoning attacks using IP fragmentation*, `https://indico.dns-oarc.net/event/31/contributions/692/attachments/660/1115/fujiwara-5.pdf`, DNS-OARC 30, (2019).
[15] Markus Brandt, Tianxiang Dai, Amit Klein, Haya Shulman, and Michael Waidner: *Domain Validation++ For MitM-Resilient PKI*, In Proceedings of the ACM SIGSAC Conference on Computer and Communications Security, pp. 2060-2076 (2018).
[16] Let's Encrypt: *EDNS Buffer Size Changing to 512 Bytes*, `https://community.letsencrypt.org/t/edns-buffer-size-changing-to-512-bytes/77945`, (2018).
[17] Internet Systems Consortium: *BIND 9 Updates, April 2017*, `https://www.isc.org/blogs/bind-april-2017/`, (2017).
[18] IETF: *Measures against cache poisoning attacks using IP fragmentation in DNS*, `https://tools.ietf.org/html/draft-fujiwara-dnsop-fragment-attack-01`, Internet-Draft, (2019).
[19] IETF: *Avoid IP fragmentation in DNS*, `https://tools.ietf.org/html/draft-fujiwara-dnsop-avoid-fragmentation-01`, Internet-Draft, (2019).