

シミュレーションによる DNSSEC の UDP トラフィック解析

力武 健次[†] 中尾 康二^{†,††} 下條 真司^{†††} 野川 裕紀^{††††}

[†] 独立行政法人 情報通信研究機構 セキュリティ高度化グループ 〒184-8795 東京都小金井市貫井北町 4-2-1

^{††} KDDI 株式会社 技術開発本部 情報セキュリティ部 〒102-0072 東京都千代田区飯田橋 3-10-10

^{†††} 大阪大学 サイバーメディアセンター 〒567-0047 大阪府茨木市美穂が丘 5-1

^{††††} 東京医科歯科大学 情報医科学センター 〒113-8510 東京都文京区湯島 1-5-45

E-mail: †rikitake@nict.go.jp

あらまし DNS(ドメイン名システム)の認証手法の1つであるDNSSECでは、DNSの応答パケットにデジタル署名を加えることでペイロード長が大きくなる。このペイロード長増大によりDNSのUDPペイロードを運ぶIPパケットの分割とパケットロス率の増加が起こり、DNSのリゾルバ-サーバ間通信の信頼性が下がると予想する。本稿では、実トラフィックのサンプルおよびDNSSECの署名を追加したペイロード長値の再計算により推定したDNSペイロード長分布のモデルを提案し、それに基づいてDNSのリゾルバ-サーバ間のパケットロスと分割の発生率をネットワークトラフィックのシミュレーションにより推定する手法について考察する。

キーワード DNS(ドメイン名システム), DNSSEC, ペイロード長, 輸送プロトコル

A Simulation-based UDP Traffic Analysis of DNSSEC

Kenji RIKITAKE[†], Koji NAKAO^{†,††}, Shinji SHIMOJO^{†††}, and Hiroki NOGAWA^{††††}

[†] Security Advancement Group, NICT, Japan 4-2-1, Nukui-Kitamachi, Koganei City, Tokyo 184-8795 Japan

^{††} Information Security Department, KDDI Corporation

3-10-10, Iidabashi, Chiyoda-Ku, Tokyo 102-0072 Japan

^{†††} Cybermedia Center, Osaka University 5-1, Mihogaoka, Ibaraki-City, Osaka 567-0047 Japan

^{††††} Information Center for Medical Sciences, Tokyo Medical and Dental University

1-5-45, Yushima, Bunkyo-Ku, Tokyo 113-8510 Japan

E-mail: †rikitake@nict.go.jp

Abstract DNSSEC, an authentication method of DNS (Domain Name System), increases the payload length of DNS answer datagrams by adding digital signatures. The payload-length increase causes fragmentation and larger loss rate of the IP datagrams which carry the DNS UDP payloads, and reduces the reliability of DNS resolver-server transactions. In this paper, we propose a model of the length distributions of DNS UDP payloads estimated from real-world traffic samples and recalculation of the payload length values after adding DNSSEC signatures. We then propose the network traffic simulation procedure to estimate the rates of loss and fragmentation of IP datagrams between DNS resolvers and servers.

Keywords DNS (Domain Name System), DNSSEC, payload length, transport protocol

1. Introduction

Authenticating DNS (Domain Name System) is a critical issue for prevent criminal activities by false identification of Internet servers, such as phishing or pharming frauds. Wide deployment of authentication infrastructure for DNS is essential, since current DNS does not have one and anyone can forge DNS answers.

DNS Security Extensions, also called as *DNSSEC*, are a collection of new resource records and protocol modifications that add data origin authentication and data integrity to the DNS, as described in the recent set of RFCs (RFC4033 [1], RFC4034 [2], and RFC4035 [3]), published on March 2005.

DNSSEC provides authentication by associating cryptographic signatures generated for each RRset (Resource Record ets) exchanged between the servers and resolvers

(clients). DNSSEC has the zone-delegation-based authentication hierarchy. DNSSEC-aware resolver must follow the authentication chain formed by pairs of DS (Delegation Signer) RR (Resource Record) and DNSKEY (DNS Public Key) RR.

Including the DNSSEC-related signature information largely increases the amount of payload of DNS answers, since DNSSEC requires all parent zones for a trusted zone must be signed. Each RRSIG (Resource Record Signature) RR includes a 128-byte signature of RSA/SHA1 [4]. We have presented an estimation result in a previous paper [5] which shows the size of DNSSEC-based DNS answers is roughly five times larger than the DNS answers without digital signatures, under the condition of at least one RRSIG RR is attached to an RRset of a DNS answer, as described in Section 2.2 of RFC4035.

Moreover, a DNS answer payload includes DNSKEY RRs of 514-byte public KSK (Key-Signing Key) and 130-byte public ZSK (Zone-Signing Key) when requested by the resolver, so the answer length could even become larger. This increase causes fragmentation of IPv4 (Internet Protocol version 4) or IPv6 datagrams carrying the UDP payloads of DNS, and may affect the reliability of DNS resolver-server transactions in UDP.

In this paper, we analyze how introduction of DNSSEC changes the UDP payload-length distribution characteristics by estimating the curves from a set of empirical data obtained from a set of real-world DNS traffic. We also propose a set of payload-length distribution models for performing network simulation to obtain the traffic characteristics.

In later sections, we first describe a historical change of DNS UDP payloads in Section 2. We then explain what may happen when UDP payload size becomes larger on DNS transactions and how the increase may affect on the overall DNS reliability in Section 3. We then propose a set of DNS traffic models for configuring and performing network simulation in Section 4. We then show the distribution curve estimation results of DNS UDP payload length in Section 5. Finally, our conclusions and future works are presented in Section 6.

2. DNS UDP Payload Length

In this section, we describe how DNS programs traditionally chooses the transport protocols based on the payload length.

DNS resolver-server exchange can use both TCP [6] and UDP [7], for maximizing the efficiency, since traditionally most of the DNS payloads are small. Keeping the whole payload in single UDP and IP datagram prevents the fragmentation of payloads and can minimize the number of exchanged datagrams for maximum efficiency. Section 6.1.3.2

Table 1 An example of usage details of EDNS0 and the specified length based on a real-world traffic sample [13].

for 3694918 answers of 28-NOV-2003 with OPT RRs (43.97% of 6249736 answers)				
UDPsize	512	1280	2048	4096
numbers	13	1399	706497	2987009
%	< 0.01	0.038	19.12	80.84
for 1318187 answers of 16-DEC-2003 with OPT RRs (59.12% of 2997881 answers)				
UDPsize	512	1280	2048	4096
numbers	1	1126	434632	882428
%	< 0.01	0.085	32.97	66.94

of RFC1123 [8] shows that a DNS server must service UDP queries and it should service TCP queries as well.

The selection between TCP and UDP is solely performed based on the length of payloads. Section 4.2.1 of RFC1035 [9] restricts the size of UDP queries and answers to 512 bytes due to the minimum UDP reassembly unit of IPv4. Larger DNS payloads must be transmitted through TCP.

Traditionally the 512-byte limitation was sufficient for DNS answers without digital signatures, according to the empirical studies of DNS payload length [10]. The RFC1035 restriction did not largely affect the DNS scalability, except for a widely-known empirical issue that the number of IPv4 Root Server addresses is limited to 13 provided that all the A RRs are included in the answer [11].

With DNSSEC, however, the length of DNS answers largely increases and most of them exceeds the 512-byte limit, due to the attached digital signature RRs. Converting the whole DNS resolver-server traffic from UDP to TCP largely increases the number of IP datagrams, and the total cost of the transactions becomes much higher.

A DNS extension called EDNS0 [12] solves this problem by allowing larger UDP payload size by mutual agreement of the requestor (resolver) and responder (cache or server). Section 3 of RFC4035 defines all “security-aware” (i.e., DNSSEC-capable) DNS servers must support the EDNS0 and a message size of at least 1220 octets, and should support a message size of 4000 octets.

Table 1 shows an example data set as of December 2003, based on the traffic data taken from Osaka University campus network. The results show that 44~59% of DNS answers included OPT RRs indicating the EDNS0 support, and that 67~81% of OPT RRs suggest server-side payload-length support of 4096 bytes (as *UDPsize*). BIND [14] 8.3.0 and later has the default value of *UDPsize* as 4096, so this suggest the popularity of BIND-based DNS servers as well as that of EDNS0.

The fact that more than 40 % of the DNS answers sup-

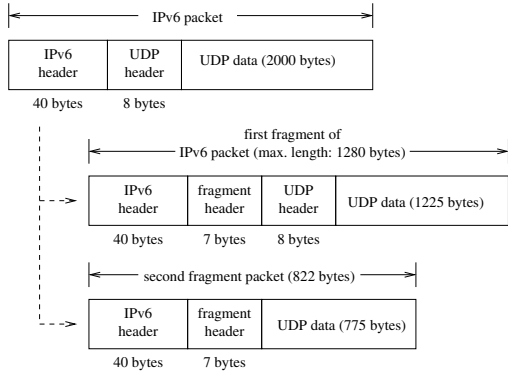


Fig. 1 An example of IPv6 UDP fragmentation [16].

ports OPT RR indicates at least the server-side support of EDNS0 is gaining popularity, although the client-side capability of EDNS0 and DNSSEC including the support of DO bit [15] is not yet under wide deployment.

3. Effects of UDP Fragmentation on DNS Reliability

In this section, we explain how larger payloads such as those of DNSSEC affects the overall DNS performance when they are transferred over the current protocol-choosing algorithm.

While enlarging the limitation of UDP payload length by EDNS0 is an effective workaround to avoid using higher-cost TCP, it also imposes another problems which may affect the DNS reliability on UDP transport.

DNS UDP transport implicitly assumes that a message between resolvers and servers is transferred by single UDP payload. While this assumption works well while UDP payloads are not fragmented by the lower layer protocols including IP, DNSSEC requirement of larger UDP datagram size breaks this assumption.

When sending a UDP datagram, the payload data and the UDP header are passed to the lower-layer IP protocol handler, as a payload of an IP datagram. Since the IP datagram has its own payload size limitation of MTU (Maximum Transmission Unit), the UDP datagram must be fragmented into multiple IP datagrams.

On IPv4, IP datagram fragmentations are automatically handled by the intermediate routers, so the sender host does not have to take a special care. On IPv6, however, the routers discard all IP datagrams exceeding the MTU, so the sender host must take care of the fragmentation.

Figure 1 shows an example of how IPv6 datagrams are fragmented by the sender. In this case, a 2000-byte UDP datagram over IPv6 datagrams are splitted into two IP datagrams of 1280 and 822 bytes. The implication of IPv6 fragmentation handling is also documented at Section 3 of RFC4035, as the servers should take steps to ensure that




Fig. 2 Delivery loss of UDP payloads fragmented into multiple IP datagrams.

UDP datagrams it transmits over IPv6 are fragmented at the minimum IPv6 MTU (1280 bytes in RFC2460 [17] Section 5) unless the path MTU is known, if necessary.

While IP fragmentation looks transparent from UDP, it may cause a degradation of reliability: a lost IP fragment causes the retransmission of the entire UDP datagram [18]. Kent and Mogul [19] claim the following three points against fragmentation:

- fragmentation causes inefficient use of resources, which can greatly increase the cost of delivering a datagram;
- loss of fragments leads to degraded performance, caused by the retransmission of all of the data in the original datagram; and
- efficient reassembly is hard, by the diversity of situations in which the reassembly process yields lower than desired performance.

Kent and Mogul [19] proposes a simple probabilistic model showing the effects of IP datagram loss for a UDP payload delivery. Assume we want to deliver a UDP payload fragmented into a set of F IP datagrams. Each of the IP datagrams has probability p of being lost or undelivered. We assume that the arrival or loss of each IP datagram is mutually independent. We use $E(N)$ to denote the expected number of datagrams that will be sent in order to successfully deliver N UDP payloads.

If any IP datagram is lost, all F fragments must be retransmitted. The probability that a datagram is successfully delivered is $1 - p$. Since each fragment arrival is independent, the probability that all the necessary datagrams are successfully delivered is $(1 - p)^F$. Thus,

$$E(N) = \frac{NF}{(1 - p)^F} \quad (1)$$

Figure 2 shows example situations of the simple probabilistic model explained above where multiple IP datagrams are required to send a UDP payload based on Equation 1. Fig. 2 shows four cases for each integer value of $F = 1$ to $F = 4$, since the current DNSSEC standard suggests sup-

porting 4000 octets for a payload, which consists of at least four IPv4 or IPv6 datagrams.

For F IP datagrams/payload, the following changes happen from the case for a UDP payload with no IP fragmentation:

- the number of IP datagrams needed to get through the payload increases F times;
- the probability of failed delivery increases to $1 - (1-p)^F$, which is $\approx 1 + Fp$ when p is very small and gets much larger when p is large;
- the latency of a UDP message delivery increases proportionally or even more as the number of IP datagrams does.

The issue of multiple-IP-datagram payloads has been addressed by Gudmundsson [20] and later by Arends et al. [3], although no quantitative analysis of how the fragmentation of IP datagram affects the overall degradation of reliability for DNS UDP payload exchanges has been given.

On DNS UDP transactions, each resolver and server must perform retries on the payload loss and determine the protocol timeout, since UDP does not support retransmission itself. The UDP retry strategies are different between the implementations. Here are some examples:

- The built-in resolver of FreeBSD [21] 4.11-RELEASE, based on BIND Version 8.3.7-REL, the minimum time between retries is set to 5 seconds, and the default number of retries is set to 2. Simply put, the queries for first DNS server is sent twice with 5-second interval^{*1}. While this simple retransmission algorithm works well on a low-latency and not-so-congested network links, this may cause congestion when all DNS resolvers behave in this way.
- On djbdns [22], the resolver routine first sends a query and waits 3 seconds for a response, then sends the same request to the other servers, waiting 3 seconds for each server for a response. If no answer is received, the resolver send a query to each server again and waits for 11 seconds for each response; if that fails again, the resolver tries again for the last time, waiting 45 seconds for each server [23]. This pseudo-exponential backoff strategy works well on a high-latency network links, the delay caused by UDP payload loss may become noticeable.

The behavior of DNS software interaction of different retransmission and timeout strategies is hard to expect algorithmically without performing simulation-based experiments. In later sections, we propose a procedure to estimate

*1 : The interval changes for each of multiple DNS servers to which the resolver sends queries. See the source code files of `lib/resolv/res_send.c` and `include/resolv.h` under `/usr/src/contrib/bind/` for the further details.

the rates of loss and fragmentations of DNS UDP payloads and the IP datagrams.

4. Simulation Model and Procedure for Large-payload DNS Traffics

The UDP payload and IP datagram behavior of DNS may drastically change by introducing DNSSEC and other larger-payload application on the Internet, and is crucial for most of the Internet application, fully dependent on the DNS name resolution service.

We are currently proceeding the evaluation of how DNSSEC affects the DNS UDP protocol behavior as follows:

- Gathering empirical information of DNS payload length:* we first collected a set of real-world traffic data [13], and obtained the empirical UDP payload length distribution as CDFs (Cumulative Distribution Functions) of QNAMEs (DNS query names) and the valid (non-error) answers.
- Building analytic models:* based on the data collected in **a)**, we build analytic models for distributions of UDP payload length sets of DNS QNAMEs in the queries and the valid answers. We will show our curve estimation results in Section 5.

Paxson [24] describes advantages of analytic models compared to empirical models as follows:

- analytic models are often mathematically tractable, lending themselves to greater understanding;
- analytic models are very concise and thus easily communicated; and
- with an analytic model, different datasets can be easily compared by comparing their corresponding estimates for the analytic model's parameters.

Giving a mathematical estimation for these two datasets helps performing a network traffic simulation.

- Building simplified DNS query-answer model for the payload length evaluation:* with the size distributions of DNS QNAMEs and the valid answers estimated by **b)**, a simplified payload-length distributions for DNS queries and answers are given as follows:

- A DNS query has a 12-byte header, QNAME, 4-byte QTYPE and QCLASS, and a 11-byte OPT RR, assuming the query complies with RFC1035, EDNS0 and RFC3225 extension rules, so the length is $(27 + (\text{QNAME size}))$ in bytes.
- A minimal error response has the same size as the DNS query, since the error code is included in the header, although the actual error response distribution is implementation-dependent^{*2}.

*2 : In actual responses to a Root Server using BIND Version 9 resolver,

- Suppose if e is the probability of returning error responses from a DNS server, the length of a DNS answer is $(27+(\text{QNAME size}))$ at probability of e , and a valid-answer length at probability of $(1 - e)$. We should note that the value of e have to be changed by the role of the server; if the server is an authoritative-only non-recursive server of a small-size DNS zone, e will be very small and close to zero. On the other hand, if the server is a Root Server, e could be as large as 0.98, according to Wessels and Fomenkov [25], who claims that *legitimate* queries to the `f.root-servers.net` on October 4, 2002, was only 2.15% of the total 152,744,325 queries. Kato and Sekiya [26] also claims only 0.2% of the answers to the Root Servers is valid, so e could be as large as 0.998. The estimation of e should be done by using empirical values.
- d) *Running network simulator with the estimated analytic models*: by using the analytic distribution models in c) and using the retransmission models obtained from the actual implementation such as the resolvers of BIND and djbdns, one can perform a network traffic simulation over a simulator such as ns [27]. The details of IP packet arrival-time distribution and other parameters are of future works.
- e) *Evaluation and feedback*: by obtaining the result of d) with various parameters, we will be able to evaluate the optimal algorithm to reduce the UDP payload length of the DNS servers while maintaining the authentication functionality of DNSSEC. Several programming workarounds are applicable, such as splitting a transaction into multiple resolver-server exchanges, or deliberately take the load off from UDP exchange to TCP, for a larger-payload traffic such as exchanging ZSKs and KSKs.

In Section 5, we focus on b) using datasets of past real-world traffics.

5. Curve Estimation of Analytic Models for DNS Packet Length Distributions

In this section, we use a previously-analyzed dataset for the curve estimation of CDFs for length values of DNS QNAMEs and the valid answers, upon which we performed a simulation for DNSSEC-signatures-added payloads [5], for the

the NXDOMAIN error response includes a SOA RR for the Root Domain. On the other hand, a different answer comes from `dnscache`, DNS recursive cache server of djbdns, without any redundant records including the OPT RR, since djbdns does not support EDNS0. The detailed analysis of the error answer distribution is of future works and out of scope of this paper.

real-world traffic data taken on December 16, 2003. We chose the following three distribution functions from various well-known distributions [28], by popularity and simplicity of CDF calculation, for attempting curve estimation using Gnuplot [29]^{*3} :

- The CDF of *Normal distribution* $F(x)$ is

$$F(x) = \frac{1}{2} \left\{ 1 + \operatorname{erf} \left(\frac{x - \mu}{\sqrt{2}\sigma} \right) \right\} \quad (2)$$

where $\operatorname{erf}(x)$ is the error function defined as

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2) dt \quad (3)$$

The mean value of the normal distribution above is μ and the standard deviation is σ .

- The CDF of *Weibull distribution* $F_w(x)$ is

$$F_w(x) = 1 - \exp \left(- \left(\frac{x}{a} \right)^b \right) \quad (4)$$

where a is the scale parameter and b is the form parameter ($a > 0, b > 0$). This distribution is commonly used in reliability engineering.

- The CDF of *Extreme-value distribution* $F_e(x)$ is

$$F_e(x) = \exp \left(- \exp \left(- \frac{x - c}{d} \right) \right) \quad (5)$$

where c is the positional parameter and d is the scale parameter ($c > 0, d > 0$). This distribution is the limiting distribution for the minimum or the maximum of a very large collection of random observations from the same arbitrary distribution, and is commonly used for estimating extreme values.

We also used the χ^2 test for measuring the discrepancy between the estimated distribution curves and the actual empirical dataset. For sets of N independent values, where the i th empirical data is E_i , and the corresponding value of a distributed curve is D_i , the X^2 value for the estimated distribution is

$$X^2 = \sum_{i=1}^N \frac{(D_i - E_i)^2}{E_i} \quad (6)$$

for the degrees of freedom is $N - 1$. For the same N , the smaller the value of X^2 , the closer the estimated distribution is fit into the empirical dataset.

We first estimated the CDF of DNS QNAMEs for each of the three distributions. Table 2 shows the statistics of the empirical dataset for the Normal distribution and the “scaled” (estimated-by-Gnuplot) Extreme-value and Weibull distributions. The X^2 values show that the Normal distribution is the best fit among the three distributions.

*3 : Gnuplot has a non-linear least-squares curve-fitting function with the Marquardt-Levenberg algorithm.

Table 2 Statistical descriptions of the estimated curves for DNS QNAME dataset of December 16, 2003.

distribution type			X^2
Normal	$\mu = 23.586$	$\sigma = 6.185$	0.1851
Weibull (scaled)	$a = 25.323$	$b = 4.520$	17.70
Extreme (scaled)	$c = 20.818$	$d = 5.115$	$> 6 \times 10^{15}$

(X^2 is the sum value of χ^2 test against the empirical CDF data for the range of $1 \leq \text{QNAME} \leq 60$, of which the *degrees of freedom* is $59 (= 60 - 1)$).




Fig. 3 Estimated distribution curves for a DNS QNAME dataset.




Fig. 4 Estimated distribution curves for a DNS valid-answer dataset without Additional Records and no signature.




Fig. 5 Estimated distribution curves for a DNS valid-answer dataset with Additional Records and no signature.

We also estimated the CDFs of DNS valid-answer lengths using another datasets obtained on the same day, of the four types of empirical data: with or without DNS Additional Records, with no signature or with simulated DNSSEC digital signatures of our previous study [5].

For this estimation, we intentionally omitted the small-

Table 3 Statistical descriptions of the estimated curves for DNS valid-answer datasets of December 16, 2003.

for 1441216 samples of December 16, 2003			
without AR and no signature			
for the range of $62 \leq \text{length} \leq 4000$ (in bytes)			
distribution type			X^2
Normal	$\mu = 136.44$	$\sigma = 43.25$	0.869
Weibull (scaled)	$a = 146.87$	$b = 3.785$	0.642
Extreme (scaled)	$c = 117.47$	$d = 33.40$	0.476
with AR and no signature			
for the range of $81 \leq \text{length} \leq 4000$ (in bytes)			
distribution type			X^2
Normal	$\mu = 197.44$	$\sigma = 76.03$	2.373
Weibull (scaled)	$a = 211.13$	$b = 3.176$	2.614
Extreme (scaled)	$c = 162.85$	$d = 55.41$	1.053
without AR and simulated DNSSEC signatures			
for the range of $262 \leq \text{length} \leq 4000$ (in bytes)			
distribution type			X^2
Normal	$\mu = 421.56$	$\sigma = 101.82$	3.195
Weibull (scaled)	$a = 451.79$	$b = 4.838$	1.920
Extreme (scaled)	$c = 376.51$	$d = 83.08$	0.378
with AR and simulated DNSSEC signatures			
for the range of $327 \leq \text{length} \leq 4000$ (in bytes)			
distribution type			X^2
Normal	$\mu = 1090.87$	$\sigma = 481.32$	40.56
Weibull (scaled)	$a = 1144.39$	$b = 3.100$	18.64
Extreme (scaled)	$c = 881.85$	$d = 303.04$	7.98

AR: Additional Records

(The lowest value of the answer length (in bytes) is the point where CDF of the value ≥ 0.01 , i.e., the data for smallest 1% of length are omitted from curve-fitting.)




Fig. 6 Estimated distribution curves for a DNS valid-answer dataset without Additional Records and with simulated DNSSEC signatures.

est 1% of length value data from the curve estimation, since DNS servers highly unlikely send shorter-length answers in the real-world practice. We chose the empirical CDF values ≥ 0.01 .

Table 3 and Figures 4, 5, 6, and 7 show the estimated and empirical distribution curves. The χ^2 test for DNS valid-length answers shows that the Extreme Value distribution fits the best among the three analytic distribution curves in




Fig.7 Estimated distribution curves for a DNS valid-answer dataset with Additional Records and with simulated DNSSEC signatures.

all four cases.

6. Conclusion and Future Works

In this paper, we described the traditional strategy of DNS programs to choose the transport protocols based on the payload length, and how the larger payloads as results of DNSSEC signature attachments affect the overall DNS performance by causing the IP-level datagram fragmentation of a UDP payload. We then proposed a procedure of evaluating the effects of large-payload DNS traffics by building analytic models from empirical datasets obtained from the real-world traffic data, using the distributions of DNS QNAME and valid-answer length. We also showed analytic model estimation for DNS QNAMEs and the valid-answers. Choosing between the three analytic distributions of Normal, Extreme-value, and Weibull distributions, we concluded that a Normal distribution was the best choice for modeling DNS QNAMEs, and a scaled Extreme-value distribution for modeling DNS valid-answer length.

For the more accurate estimation of the analytic model of DNS traffic, we need to obtain more diversified types of dataset such as DNS error-or-invalid answers and the probabilistic model of DNS answer response between valid and invalid answers, as well as datasets from more diversified sources such as those from large-scale DNS caches of ISPs (Internet Service Providers), though we need to minimize the quantity and number of attributes of the datasets from ISPs to prevent legal issues caused by monitoring the contents of DNS payloads.

We also need to perform actual network traffic simulation activities using the analytic models of this paper, as well as determining many other parameters such as time intervals of datagram arrival to the DNS server and the implementation of DNS resolver retransmission algorithms.

Acknowledgements

We thank ODINS (Osaka Daigaku Information Network System) staff members and KDDI R&D Laboratories, Inc., for allowing us to reuse the derivatives of monitored DNS payload data for this research, as well as supporting our past research activities on DNS security and performance.

Our thanks also go to Mr. Tohru Asami, the president and CEO of KDDI R&D Laboratories, Inc., Mr. Kazunori Fujiwara of Japan Registry Services, Co., Ltd., and Prof. Youki Kadobayashi of Nara Advanced Institute for Science and Technology, for supporting our research activities and valuable suggestions on DNS protocol and security issues.

References

- [1] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS security introduction and requirements," March 2005. RFC4033.
- [2] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "Resource records for the DNS security extensions," March 2005. RFC4034.
- [3] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "Protocol modifications for the DNS security extensions," March 2005. RFC4035.
- [4] D. Eastlake, "RSA/SHA-1 SIGs and RSA KEYs in the Domain Name System (DNS)," May 2001. RFC3110.
- [5] K. Rikitake, H. Nogawa, T. Tanaka, K. Nakao, and S. Shimojo, "An analysis of DNSSEC transport overhead increase," IPSJ SIG Technical Reports 2005-CSE-28, vol.2005, no.33, pp.345–350, March 2005. ISSN 0919-6072.
- [6] J. Postel, "Transmission control protocol," 1981. RFC793 (also STD7).
- [7] J. Postel, "User datagram protocol," 1980. RFC768 (also STD6).
- [8] R. Braden (Editor), "Requirements for Internet hosts – application and support," 1989. RFC1123.
- [9] P.V. Mockapetris, "Domain names – implementation and specification," 1987. RFC1035 (also STD13).
- [10] K. Rikitake, H. Nogawa, T. Tanaka, K. Nakao, and S. Shimojo, "DNS Transport Size Issues in IPv6 Environment," Proceedings of the 2004 International Symposium of Applications and the Internet (SAINT2004) Workshops, pp.141–145, 2004. ISBN 0-7695-2050-2/04.
- [11] P. Vixie and A. Kato, "DNS Response Size Issues," July 2004. INTERNET-DRAFT draft-ietf-dnsop-resize-01.txt.
- [12] P. Vixie, "Extension Mechanisms for DNS (EDNS0)," 1999. RFC2671.
- [13] K. Rikitake, H. Nogawa, T. Tanaka, K. Nakao, and S. Shimojo, "An Analysis of DNS Payload Length Increase during Transition to IPv6," IEICE Trans. Commun. (Japanese Edition), vol.J87-B, no.10, pp.1552–1563, Oct. 2004.
- [14] Internet Software Consortium, "BIND." <http://www.isc.org/bind/>.
- [15] D. Conrad, "Indicating resolver support of DNSSEC," Dec. 2001. RFC3225.
- [16] K. Rikitake, A Study of DNS Transport Protocol for Improving The Reliability, Ph.D. dissertation, Graduate School of Information Science and Technology, Osaka University, Osaka, Japan, Dec. 2004. Public Release Version 1.0, <http://www.ne.jp/asahi/bdx/info/depot/dnstransport-phd-v10-pub.pdf>.

- [17] S. Deering and R. Hinden, “Internet Protocol, Version 6 (IPv6) specification,” 1998. RFC2460.
- [18] W.R. Stevens, TCP/IP Illustrated, Volume 1, Addison–Wesley, 1994.
- [19] C.A. Kent and J.C. Mogul, “Fragmentation considered harmful,” 1987. Digital Equipment Corporation Western Research Laboratory Technical Report 87.3, <http://research.compaq.com/wrl/techreports/abstracts/87.3.html>.
- [20] O. Gudmundsson, “DNSSEC and IPv6 A6 aware server/resolver message size requirements,” Dec. 2001. RFC3226.
- [21] The FreeBSD Project, “FreeBSD.” <http://www.freebsd.org/>.
- [22] D.J. Bernstein, “djbdns.” <http://cr.yp.to/djbdns.html>.
- [23] D.J. Bernstein, “User’s guide to name resolution.” <http://cr.yp.to/djbdns/resolve.html>.
- [24] V. Paxson, “Empirically derived analytic models of wide-area TCP connections,” IEEE/ACM Trans. Netw., vol.2, no.4, pp.316–336, 1994.
- [25] D. Wessels and M. Fomenkov, “Wow, That’s a Lot of Packets,” Proceedings of the Passive and Active Measurement Workshop 2003 (PAM2003), NLANR/MNA, April 2003.
- [26] A. Kato and Y. Sekiya, “Analysis of DNS Traffic at a DNS Server in an ISP,” IEICE Trans. Commun. (Japanese Edition), vol.J87-B, no.3, pp.327–335, March 2004.
- [27] USC/ISI, “The Network Simulator - ns-2.” <http://www.isi.edu/nsnam/ns/>.
- [28] S. Nishizawa, A Probabilistic and Statistical Theory and Numerical Experiment on Trends of Meteorological Changes, Ph.D. dissertation, Graduate School of Science, Kyoto University, Kyoyo, Japan, 2005. <http://www-mete.kugi.kyoto-u.ac.jp/seiya/pdfs/nishizawa-thesis.pdf>.
- [29] T. Williams and C. Kelley, “Gnuplot.” <http://www.gnuplot.info/>.