社団法人　情報処理学会　研究報告
IPSJ SIG Technical Report

2005－DPS－122（60）
2005－CSEC－28（60）
2005／3／23

DNSSEC トランスポートオーバヘッド増加に関する解析

力武 健次 ◇‡　野川 裕記 *　田中 俊昭 ‡　中尾 康二 †　下條 真司 ⋆

DNSSEC は現在 DNS（ドメイン名システム）認証の事実上の標準となるべく再設計の段階にある．DNSSEC ではすべての RRset（リソースレコード集合）にデジタル署名を付けなければならないため，UDP トランスポートでのペイロード長が大きく増える．本論文では，実トラフィックのサンプルに対する DNSSEC プロトコルに基づいたシミュレーションを行うことで， DNSSEC の署名および関連する RR（リソースレコード）が引き起こすペイロード長の増加の影響について解析する．シミュレーションの結果は，additional records を含む DNS 応答のペイロード長のうち，IPv6 の MTU（最大送信単位）による現実的な制約である 1232 バイトを越えるものは，サンプル全体のほぼ 30%に達することを示している．

An Analysis of DNSSEC Transport Overhead Increase

Kenji Rikitake◇‡, Hiroki Nogawa*, Toshiaki Tanaka‡, Koji Nakao†, and Shinji Shimojo⋆

DNSSEC has been under the redesign process to become a de-facto standard of DNS (Domain Name System) authentication. DNSSEC requires a digital signature to be attached to every RRset (Resource Record set) of the answers, and largely increase the payload length of the UDP transport. In this paper, we analyze the effect of payload-length increase caused by DNSSEC signature and related RRs (Resource Records) by performing a simulation based on the DNSSEC protocol over real-world traffic samples. The simulation results indicate the percentage of payload length of the DNS answers including additional records exceeding the practical limitation of 1232 bytes imposed by IPv6 default MTU (Maximum Transmission Unit) becomes approx. 30% of the samples.

1 Introduction

DNS (Domain Name System) authentication has become a critical issue to prevent criminal activities using fabricated DNS names and DNS server hijacking, such as phishing fraud. DNSSEC [1] is an authentication extension of DNS, which has been under development since 1997 proposed first as RFC2065, and now under redesign process introducing DS (Delegation Signer) [2] RR (Resource Record), which intends to reduce complicated public-key signing issues between parent and child zones.

The feasibility study of the current DNSSEC proposal is still ongoing and some vulnerabilities are found, such as that of revealing all zone RRs by tracing NSEC RRs [3]. Since DNSSEC requires all answers to be digitally signed, the payload length will become much larger. While RFC3226 [4] Section 3 suggests that all DNSSEC-capable resolvers must support EDNS0 [5] to handle extended payload length larger than the current limit of 512 bytes imposed by RFC1035 [6]

◇ 大阪大学 大学院 情報科学研究科 / Graduate School of Information Science and Technology, Osaka University
* 東京医科歯科大学 情報医科学センター / Information Center for Medical Sciences, Tokyo Medical and Dental University
‡ 株式会社 KDDI 研究所 セキュリティグループ / Security Laboratory, KDDI R&D Laboratories, Inc.
† KDDI 株式会社 技術開発本部 情報セキュリティ技術部 / Information Security Department, KDDI Corporation
⋆ 大阪大学 サイバーメディアセンター / Cybermedia Center, Osaka University

```
;; QUESTION SECTION:
;abc.local. IN MX

;; ANSWER SECTION:
abc.local. 300 IN MX 0 abc.local.

;; AUTHORITY SECTION:
local. 300 IN NS ns1.local.
local. 300 IN NS ns2.local.

;; ADDITIONAL SECTION:
abc.local. 300 IN A 192.168.1.3
abc.local. 300 IN A 192.168.1.1
abc.local. 300 IN A 192.168.1.2
ns1.local. 300 IN A 127.1.0.1
ns2.local. 300 IN A 127.1.0.2
```

Fig. 1 An example of non-DNSSEC DNS answer

Section 4.2.1, very few detailed analysis has been made to actual payload increase and overhead to be imposed by DNSSEC.

We have proven that there will be a significant increase on DNS transport overhead during IPv6 migration phase [7, 8], by performing simulation by recalculating UDP payload length of each DNS answers with a scenario-based model and a set of collected real-world DNS traffic. The same approach is also applicable to predict the transport overhead by the payload length increase caused by other protocol extensions, including DNSSEC.

In this paper, we perform a simulation with a scenario and algorithm derived from the current DNSSEC implementation available as BIND [9] Version 9.3.0 [*1], and based on the proposed Internet-Drafts on DNSSEC protocol modifications [11], on DNSSEC RRs [12], NSEC RR [13] and the RSA/SHA1 [14] signature and public key algorithm, with a set of real-world DNS traffic.

In later sections, we describe how DNSSEC works and how we calculate a payload length of DNSSEC-signed answers in Section 2. We then show the simulation results in Section 3, and the conclusions and future works in Section 4.

[*1] ISC (Internet Software Consortium) suggests to upgrade to Version 9.3.1 to work around a validator vulnerability [10] of DNSSEC. The vulnerability does not affect the issues described in this paper.

```
;; QUESTION SECTION:
;abc.local. IN MX

;; ANSWER SECTION:
abc.local. 300 IN MX 0 abc.local.
abc.local. 300 IN RRSIG MX 5 4 300
                  20050316121816 20050214121816
                  42674 local.
                  (Base64 128-byte signature)

;; AUTHORITY SECTION:
local. 300 IN NS ns2.local.
local. 300 IN NS ns1.local.
local. 300 IN RRSIG NS 5 3 300
                    20050316121816 20050214121816
                    42674 local.
                    (Base64 128-byte signature)

;; ADDITIONAL SECTION:
abc.local. 300 IN A 192.168.1.2
abc.local. 300 IN A 192.168.1.3
abc.local. 300 IN A 192.168.1.1
ns1.local. 300 IN A 127.1.0.1
ns2.local. 300 IN A 127.1.0.2
abc.local. 300 IN RRSIG A 5 4 300
                  20050316121816 20050214121816
                  42674 local.
                  (Base64 128-byte signature)
ns1.local. 300 IN RRSIG A 5 4 300
                  20050316121816 20050214121816
                  42674 local.
                  (Base64 128-byte signature)
ns2.local. 300 IN RRSIG A 5 4 300
                  20050316121816 20050214121816
                  42674 local.
                  (Base64 128-byte signature)
;;; the following DNSKEY RRs are optional
;;; for KSK (Key-Signing Key, 4096 bit)
local. 300 IN DNSKEY 256 3 5
              (Base64 514-byte public-key field)
;;; for ZSK (Zone-Signing Key, 1024 bit)
local. 300 IN DNSKEY 256 3 5
              (Base64 130-byte public-key field)
```

Fig. 2 An example of DNSSEC-signed DNS answer

2 Simulation Model and Method

Figure 1 shows a typical DNS query and answer pair *without* DNSSEC. RFC2181 [15] Section 5 defines that an RRset is a group of resource records with the same label, class and type, but with different data. To secure the contents of each RRs, DNSSEC mandates signing each RRset as shown in Figure 2, which is a modified excerpt from actual output of BIND Version 9.3.0.

Current DNSSEC protocol mandates that each RRset in the Answer Section and Authority Section of a DNS answer, as defined in RFC1034 [16] Section 3.7, must be signed by at least one RRSIG

L_m = total payload length without AR
L_t = total payload length with AR

$$A_m = \sum_{\text{RRsets in AS}} (\text{Length of RRSIG}) + \sum_{\text{RRsets in NS}} (\text{Length of RRSIG})$$

// in case when only delegated NS RRs are answered
if ((no RR in AS) and (any RR exists in NS)) {
    A_{ds} = max((Length of DS), (Length of NSEC))
} else {
    $A_{ds} = 0$
}

L_{md} = total length of DNSSEC-signed payload without AR
$= L_m + A_m + A_{ds}$

L_{td} = total length of DNSSEC-signed payload with AR
$= L_t + A_m + A_{ds} + \sum_{\text{RRsets in AR}} (\text{Length of RRSIG})$

AS: RRs of Answer Section
NS: RRs of Authority Section
AR: RRs of Additional Section

- domain names are *case-insensitive*, so as used in determining the RRsets among answered RRs
- *zonename* of a domain name is defined as the right-hand-side part of the domain name whose leftmost label is stripped (e.g., `example.com` is the zonename for `host.example.com`)
- Length of RRSIG, DS, NSEC depend on the signer's zonenames referred in each RR
- For RRSIG, the signer name is the zonename of the referred name of each RRset
- For DS and NSEC, the signer name is the zonename of the queried name in the QNAME field

Fig. 3 Simulation algorithm for calculating the length of DNSSEC-signed payloads

RR. And for a reference to a delegated zone, such as in a case when no RR in the Answer Section and one or more RRs referring to an authoritative server using NS RR in the Authority Section, a DS or an NSEC RR and the corresponding RRSIG RR must be returned.

Figure 3 shows the calculation algorithm of the simulation. We assume that the DNSSEC-signed payloads contains at least one RRSIG RR for every necessary RRset to be signed. RRSIG RR length values are calculated based on actual zonenames of the RRsets.

We also assume inclusion of the DS or NSEC RR, in case when only delegated NS RRs are answered. While the detailed behavior of including DS and NSEC RR are dependent on the parent-child relationship of zones served by each DNS server and cannot be easily computed, we assume that only one of DS or NSEC RR is returned, since in most of the cases the RRs in the Authority Section contains one RRset. We assume to include larger length value of DS or NSEC RR to make the simulation results deterministic and reproducible.

By the algorithm described in Figure 3, the DNSSEC-signed length values (L_{md} and L_{td}) can be calculated from given traffic data, by the payload length, DNS payload header, and the included RRset information.

We also assume the DNSSEC length constants as:

- the length of zone-signing key is 1024 bits (or 128 bytes);
- the signature algorithm is RSA/SHA1; and
- the length of signature in RRSIG is 1024 bits (or 128 bytes), based on the actual measurement of DNSSEC packets.

Table 1 Statistics for the samples and the simulation results

for 1441216 samples of 16-DEC-2003			
	μ	σ	max
raw w/o AR (L_m)	136.44	43.25	944
raw with AR (L_t)	197.44	76.03	987
signed w/o AR (L_{md})	421.56	101.82	1332
signed with AR (L_{td})	1090.87	481.32	3730

μ: mean value (bytes)
σ: unbiased standard deviation (bytes)
max: maximum payload length (bytes)
w/o AR: without Additional Records
signed: simulated DNSSEC-signatures added

3 Simulation Results and Evaluation

For the simulation, we used a set of real-world traffic data collected from Osaka University campus network, of DNS UDP packets between inside and outside of the network. We filtered out malformed packets, including UDP packets of non-

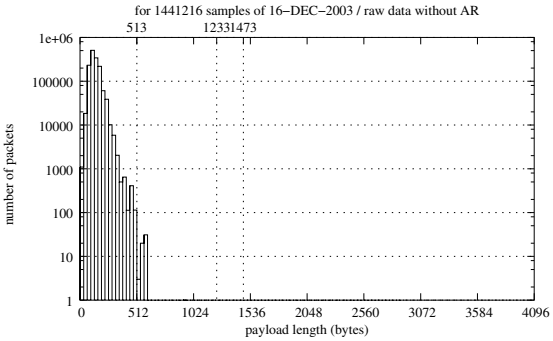

Fig. 4 Payload length of DNS answers without additional records (L_m) of 1441216 samples collected on 16-DEC-2003

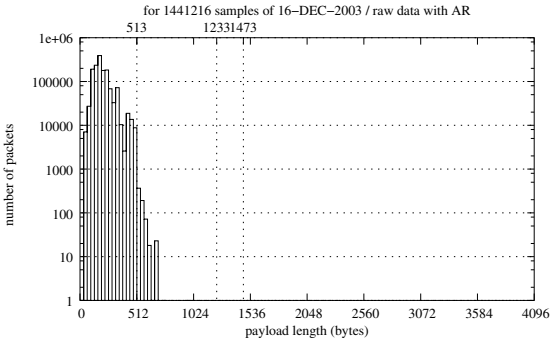

Fig. 5 Payload length of DNS answers with additional records (L_t) of 1441216 samples collected on 16-DEC-2003

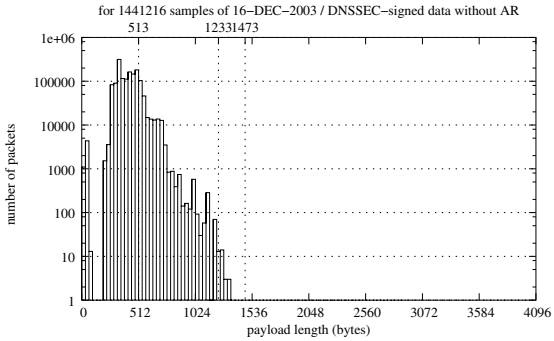

Fig. 6 Payload length of DNS answers including simulated DNSSEC signatures without additional records (L_{md}) of 1441216 samples collected on 16-DEC-2003

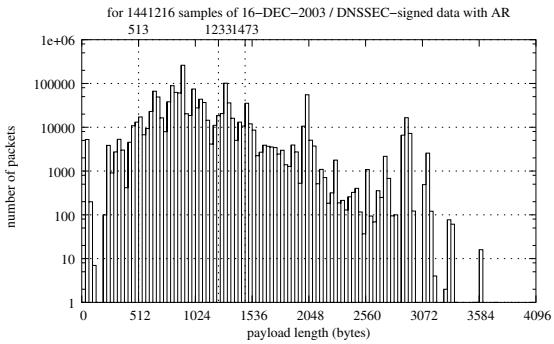

Fig. 7 Payload length of DNS answers including simulated DNSSEC signatures with additional records (L_{td}) of 1441216 samples collected on 16-DEC-2003

zero fragment offsets and those with malformed QNAMEs (such as those which includes whitespace character), and obtained 1441216 samples during 0047∼1246JST (Japan Standard Time) of December 16, 2003.

Table 1 shows the statistics of the collected traffic data and the simulation results. The mean value of L_{md} is ≃3 times larger than that of L_m, and that of L_{td} is ≃5.5 times larger than that of L_t. While all values sampled and simulated are within the limit of 4000 bytes as suggested in RFC3226, this increase will cause significant bandwidth increase of DNS UDP traffic.

Figures 4, 5, 6, and 7 show the distributions of the number of packets for payload length, of L_m, L_t, L_{md}, L_{td}, respectively. Comparing Figures 5 and 7 indicates that the numbers of packet for larger payload lengths are still high even if the length value exceeds 1220 bytes, which is the required minimum *MUST* value of the payload length in RFC3226 and the updated DNSSEC protocol draft documents.

We define our own criteria to evaluate the payload length by whether it exceeds the following three threshold values (in bytes): 512, 1232, 1472. 512 bytes is the limit of UDP payload for non-DNSSEC DNS programs. 1232 bytes is the limit of UDP payload length in an IPv6 packet, whose MTU (Maximum Transmission Unit) is 1280. 1472 bytes is the limit of UDP payload length in an IPv4 packet, whose MTU is 1500.

Table 2 shows the values of CDFs (Cumulative Distribution Functions) for the samples and the simulated results. While those values indicate the condition $L_{md} \leq 1232$ is always true so the DNSSEC-signed payloads *without Additional*

Table 2 CDFs of payload length for the samples and the simulation results

for 1441216 samples of 16-DEC-2003			
	>512	>1232	>1472
raw w/o AR (L_m)	1.0000	1.0000	1.0000
raw with AR (L_t)	0.9995	1.0000	1.0000
signed w/o AR (L_{md})	0.8473	1.0000	1.0000
signed with AR (L_{td})	0.0350	0.7064	0.8548

>512: >512 bytes (non-DNSSEC DNS limit)
>1232: >1232 bytes (IPv6 MTU limit)
>1472: >1472 bytes (IPv4 MTU limit)

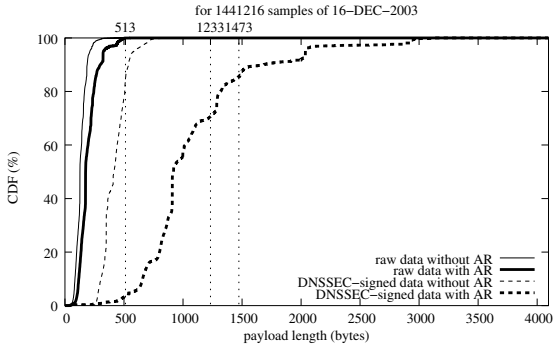

Fig. 8 The CDFs of DNS answers and results of simulation including DNSSEC signatures for 1441216 samples collected on 16-DEC-2003

Section RR, we still consider that ≃29.4% of the payloads does not fit in single unfragmented IPv6 UDP packet, and ≃14.5% of the payloads does not fit into single unfragmented IPv4 UDP packet either. This indicates wide-deployment of DNSSEC will result in a large number of fragmented UDP packets for DNS answers, and that the fragmentation will impact the reliability of UDP transport of the DNSSEC.

4 Conclusion and Future Works

In this paper, we analyzed the effect of payload-length increase caused by DNSSEC signature and related RRs by performing a simulation based on the DNSSEC protocol over real-world traffic samples from a campus network. The simulation results indicated that the percentage of payload length of the DNS answers including additional records exceeding the practical limitation of 1232 bytes imposed by IPv6 default MTU becomes ≃30% of the samples.

RRs in Additional Section (ARs) on DNS answers play significant roles on DNS lookups, by giving hints to the resolvers so that they do not have to perform redundant lookups. In some implementations, complete removal of the ARs may cause malfunction on DNS lookups [17]. While selectively responding partial RRset for IP-address RRs is effective for reduction of the payload length, the room of choice for the DNSSEC protocol is largely reduced by the requirement of sending at least one RRSIG RR for each section in DNS answer. To avoid UDP fragmentation on DNSSEC, the servers should carefully choose the strategy to trim the output RRs in the answer payload, though eliminating the fragmentation will not be possible when guaranteeing a sufficient number of ARs in DNS answer payloads is required.

We consider that the payload-length increase which is caused by IPv6 deployment should be included into the future evaluation of DNSSEC, due to the increase of the length of IP-address RR by changing from A RRs to AAAA RRs. During the IPv4-to-IPv6 migration, the servers will require to send *two* RRSIGs in a payload, each for A RRsets and AAAA RRsets, so the impact to the overall transport overhead will become even larger.

Acknowledgements

Our thanks go to Mr. Tohru Asami, the president and CEO of KDDI R&D Laboratories, Inc., Prof. Motonori Nakamura of Kyoto University, Prof. Youki Kadobayashi of Nara Advanced Institute for Science and Technology, and Prof. Akira Kato of University of Tokyo, for supporting our research activities and valuable suggestions on DNS protocol and security issues.

We also thank ODINS (Osaka Daigaku Information Network System) staff members for their kind assistance to provide a traffic monitoring environment.

References

[1] Eastlake, D.: Domain Name System Security Extensions (1999). RFC2535.

[2] Gudmundsson, O.: Delegation Signer (DS) Resource Record (RR) (2003). RFC3658.

[3] Mori, K., Fujiwara, K. and Jinmei, T.: Discussion on whether DNSSEC is required or not. distributed by WIDE Project, WIDE draft wide-draft-dns-dnssec-deployment-discussion-01.txt.

[4] Gudmundsson, O.: DNSSEC and IPv6 A6 aware server/resolver message size requirements (2001). RFC3226.

[5] Vixie, P.: Extension Mechanisms for DNS (EDNS0) (1999). RFC2671.

[6] Mockapetris, P. V.: Domain names – implementation and specification (1987). RFC1035 (also STD13).

[7] Rikitake, K., Nogawa, H., Tanaka, T., Nakao, K. and Shimojo, S.: 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.

[8] Rikitake, K., Nogawa, H., Tanaka, T., Nakao, K. and Shimojo, S.: An Analysis of DNS Payload Length Increase during Transition to IPv6, *IEICE Trans. Commun. (Japanese Edition)*, Vol. J87-B, No. 10, pp. 1552–1563 (2004).

[9] Internet Software Consortium: BIND. `http://www.isc.org/bind/`.

[10] US-CERT: BIND 9.3.0 vulnerable to denial of service in validator code. US-CERT Vulnerability Note VU#938617, `http://www.kb.cert.org/vuls/id/938617`.

[11] Arends, R., Austein, R., Larson, M., Massey, D. and Rose, S.: Protocol Modifications for the DNS Security Extensions (2004). INTERNET-DRAFT draft-ietf-dnsext-dnssec-protocol-09.txt.

[12] Arends, R., Austein, R., Larson, M., Massey, D. and Rose, S.: Resource Records for the DNS Security Extensions (2004). INTERNET-DRAFT draft-ietf-dnsext-dnssec-records-11.txt.

[13] J. Schlyter (Ed.): DNS Security (DNSSEC) NextSECure (NSEC) RDATA Format (2004). RFC3845.

[14] Eastlake, D.: RSA/SHA-1 SIGs and RSA KEYs in the Domain Name System (DNS) (2001). RFC3110.

[15] Elz, R. and Bush, R.: Clarification to the DNS Specification (1997). RFC2181.

[16] Mockapetris, P. V.: Domain names – concepts and facilities (1987). RFC1034 (also STD13).

[17] Japan Registry Service (JPRS): On Maximum Number of DNS Servers (2003). `http://jprs.jp/tech/jp-dns-info/2003-07-10-max-number-of-dns-server.html` (written in Japanese).