

Internet Engineering Task Force (IETF)
Request for Comments: 8467
Category: Experimental
ISSN: 2070-1721

A. Mayrhofer
nic.at GmbH
October 2018

Padding Policies for Extension Mechanisms for DNS (EDNS(0))

Abstract

RFC 7830 specifies the "Padding" option for Extension Mechanisms for DNS (EDNS(0)) but does not specify the actual padding length for specific applications. This memo lists the possible options ("padding policies"), discusses the implications of each option, and provides a recommended (experimental) option.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8467>.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Terminology	2
3. General Guidance	3
4. Padding Strategies	3
4.1. Recommended Strategy: Block-Length Padding	3
4.2. Other Strategies	5
4.2.1. Maximal-Length Padding	5
4.2.2. Random-Length Padding	5
4.2.3. Random-Block-Length Padding	6
5. IANA Considerations	6
6. Security Considerations	6
7. References	7
7.1. Normative References	7
7.2. Informative References	7
Appendix A. Padding Policies That Are Not Sensible	8
A.1. No Padding	8
A.2. Fixed-Length Padding	8
Acknowledgements	9
Author's Address	9

1. Introduction

[RFC7830] specifies the Extension Mechanisms for DNS (EDNS(0)) "Padding" option, which allows DNS clients and servers to artificially increase the size of a DNS message by a variable number of bytes, hampering size-based correlation of encrypted DNS messages.

However, RFC 7830 deliberately does not specify the actual length of padding to be used. This memo discusses options regarding the actual size of padding, lists advantages and disadvantages of each of these "padding strategies", and provides a recommended (experimental) strategy.

Padding DNS messages is useful only when transport is encrypted using protocols such as DNS over Transport Layer Security [RFC7858], DNS over Datagram Transport Layer Security [RFC8094], or other encrypted DNS transports specified in the future.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. General Guidance

EDNS(0) options space: The maximum message length, as dictated by the protocol, limits the space for EDNS(0) options. Since padding will reduce the message space available to other EDNS(0) options, the "Padding" option MUST be the last EDNS(0) option applied before a DNS message is sent.

Resource Conservation: Especially in situations where networking and processing resources are scarce (e.g., battery-powered long-life devices, low bandwidth, or high-cost links), the trade-off between increased size of padded DNS messages and the corresponding gain in confidentiality must be carefully considered.

Transport Protocol Independence: The message size used as input to the various padding strategies MUST be calculated excluding the potential extra 2-octet length field used in TCP transport. Otherwise, the padded (observable) size of the DNS packets could significantly change between different transport protocols and reveal an indication of the original (unpadded) length. For example, given a Block-Length Padding strategy with a block length of 32 octets and a DNS message with a size of 59 octets, the message would be padded to 64 octets when transported over UDP. If that same message were transported over TCP and the padding strategy considered the extra 2 octets of the length field (61 octets in total), the padded message would be 96 octets long (as the minimum length of the "Padding" option is 4 octets).

4. Padding Strategies

This section contains a recommended strategy, as well as a non-exhaustive list of other sensible strategies, for choosing padding length. Note that, for completeness, Appendix A contains two more strategies that are not sensible.

4.1. Recommended Strategy: Block-Length Padding

Based on empirical research performed by Daniel K. Gillmor [NDSS-PADDING], padding SHOULD be performed following the Block-Length Padding strategy as follows:

- (1) Clients SHOULD pad queries to the closest multiple of 128 octets.
- (2) If a server receives a query that includes the EDNS(0) "Padding" option, it MUST pad the corresponding response (see Section 4 of RFC 7830) and SHOULD pad the corresponding response to a multiple of 468 octets (see below).

Note that the recommendation above only applies if the DNS transport is encrypted (see Section 6 of RFC 7830).

In Block-Length Padding, a sender pads each message so that its padded length is a multiple of a chosen block length. This creates a greatly reduced variety of message lengths. An implementor needs to consider that even the zero-length "Padding" option increases the length of the packet by 4 octets.

Options: Block length. For queries, values between 16 and 128 octets were discussed before empiric research was performed. Responses will require larger block sizes (see [NDSS-PADDING] and above for a discussion).

Very large block lengths will have confidentiality properties similar to the Maximal-Length Padding strategy (Section 4.2.1), since almost all messages will fit into a single block. Such "very large block length" values are:

- o 288 bytes for the query (the maximum size of a one-question query over TCP, without any EDNS(0) options) and
- o the EDNS(0) buffer size of the server for the responses.

Advantages: This policy is reasonably easy to implement, reduces the variety of message ("fingerprint") sizes significantly, and does not require a source of (pseudo) random numbers, since the padding length required can be derived from the actual (unpadded) message.

Disadvantage: Given an unpadded message and the block size of the padding (which is assumed to be public knowledge once a server is reachable), the size range of a padded message can be predicted. Therefore, the minimum length of the unpadded message can be inferred.

The empirical research cited above performed a simulation of padding, based on real-world DNS traffic captured on busy recursive resolvers of a research network. The evaluation of the performance of individual padding policies was based on a "cost to attacker" and "cost to defender" function, where the "cost to attacker" was defined as the percentage of query/response pairs falling into the same size bucket and "cost to defender" was defined as the size factor between padded and unpadded messages. Padding with a block size of 128 bytes on the query side and 468 bytes on the response side was considered the optimum trade-off between defender and attacker cost. The response block size of 468 was chosen so that 3 blocks of 468 octets would still comfortably fit into typical Maximum Transmission Unit (MTU) size values.

The block size will interact with the MTU size. Especially for length values that are a large fraction of the MTU, unless the block length is chosen so that a multiple just fits into the MTU, Block-Length Padding may cause unnecessary fragmentation for UDP-based delivery. Of course, choosing a block length larger than the MTU always forces fragmentation.

Note: Once DNSSEC-validating clients become more prevalent, observed size patterns are expected to change significantly. In that case, the recommended strategy might need to be revisited.

4.2. Other Strategies

4.2.1. Maximal-Length Padding

In Maximal-Length Padding, the sender pads every message to the maximum size allowed by protocol negotiations.

Advantages: Maximal-Length Padding, when combined with encrypted transport, provides the highest possible level of message-size confidentiality.

Disadvantages: Maximal-Length Padding is wasteful and requires resources on the client, all intervening networks and equipment, and the server. Depending on the negotiated size, this strategy will commonly exceed the MTU and result in a consistent number of fragments, reducing delivery probability when datagram-based transport (such as UDP) is used.

Due to resource consumption, Maximal-Length Padding is NOT RECOMMENDED.

4.2.2. Random-Length Padding

When using Random-Length Padding, a sender pads each message with a random amount of padding. Due to the size of the "Padding" option itself, each message size is increased by at least 4 octets. The upper limit for padding is the maximum message size. However, a client or server may choose to impose a lower maximum padding length.

Options: Maximum and minimum padding length.

Advantages: Theoretically, this policy should create a natural distribution of message sizes.

Disadvantage: Random-Length Padding allows an attacker who can observe a large number of requests to infer the length of the original value by observing the distribution of total lengths.

According to the limited empirical data available, Random-Length Padding exposes slightly more entropy to an attacker than Block-Length Padding. Because of that, and the risk outlined above, Random-Length Padding is NOT RECOMMENDED.

4.2.3. Random-Block-Length Padding

This policy combines Block-Length Padding with a random component. Specifically, a sender randomly chooses between a few block length values and then applies Block-Length Padding based on the chosen block length. The random selection of block length might even be reasonably based on a "weak" source of randomness, such as the transaction ID of the message.

Options: Number of and the values for the set of block lengths;
source of randomness

Advantages: Compared to Block-Length Padding, this creates more variety in the resulting message sizes for a certain individual original message length.

Disadvantage: Requires more implementation effort compared to simple Block-Length Padding.

Random-Block-Length Padding requires further empirical study, as do other combinations of padding strategies.

5. IANA Considerations

This document has no IANA actions.

6. Security Considerations

The choice of the right padding policy (and the right parameters for the chosen policy) has a significant impact on the resilience of encrypted DNS against size-based correlation attacks. Therefore, any implementor of the "Padding" option must carefully consider which policies to implement, the default policy chosen, which parameters to make configurable, and the default parameter values.

No matter how carefully a client selects their padding policy, this effort can be jeopardized if the server chooses to apply an ineffective padding policy to the corresponding response packets. Therefore, a client applying the "Padding" option may want to choose a DNS server that applies a padding policy on responses that is at least equally effective.

Note that even with encryption and padding, it might be trivial to identify that the observed traffic is DNS. Also, padding does not prevent information leaks via other side channels (particularly timing information and number of query/response pairs). Countermeasures against such side channels could include injecting artificial "cover traffic" into the stream of DNS messages or delaying DNS responses by a certain amount of jitter. Such strategies are out of the scope of this document. Additionally, there is not enough theoretic analysis or experimental data available to recommend any such countermeasures.

7. References

7.1. Normative References

[NDSS-PADDING]

Gillmor, D., "Empirical DNS Padding Policy", March 2017, <<https://dns.cmrq.net/ndss2017-dprive-empirical-DNS-traffic-size.pdf>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC7830] Mayrhofer, A., "The EDNS(0) Padding Option", RFC 7830, DOI 10.17487/RFC7830, May 2016, <<https://www.rfc-editor.org/info/rfc7830>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

7.2. Informative References

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/info/rfc7858>>.

[RFC8094] Reddy, T., Wing, D., and P. Patil, "DNS over Datagram Transport Layer Security (DTLS)", RFC 8094, DOI 10.17487/RFC8094, February 2017, <<https://www.rfc-editor.org/info/rfc8094>>.

Appendix A. Padding Policies That Are Not Sensible

A.1. No Padding

In the No Padding policy, the "Padding" option is not used, and the size of the final (actually, "non-padded") message obviously exactly matches the size of the unpadded message. Even though this "non-policy" seems redundant in this list, its properties must be considered for cases in which just one of the parties (client or server) applies padding.

Also, this policy is required when the remaining message size of the unpadded message does not allow for the "Padding" option to be included -- i.e., there are fewer than 4 octets left.

Advantages: This policy requires no additional resources on the client, server, and network side.

Disadvantages: The original size of the message remains unchanged; hence, this approach provides no additional confidentiality.

The No Padding policy MUST NOT be used unless message size disallows the use of the "Padding" option.

A.2. Fixed-Length Padding

In Fixed-Length Padding, a sender chooses to pad each message with a padding of constant length.

Options: Actual length of padding

Advantages: Since the padding is constant in length, this policy is very easy to implement and at least ensures that the message length diverges from the length of the original packet (even if only by a fixed value).

Disadvantage: Obviously, the amount of padding is easily discoverable from a single unencrypted message or by observing message patterns. When a public DNS server applies this policy, the length of the padding hence must be assumed to be public knowledge. Therefore, this policy is (almost) as useless as the No Padding policy described above.

The Fixed-Length Padding policy MUST NOT be used except for test applications.

Acknowledgements

Daniel K. Gillmor performed empirical research out of which the "Recommended Strategy" was copied. Stephane Bortzmeyer and Hugo Connerly provided text. Shane Kerr, Sara Dickinson, Paul Hoffman, Magnus Westerlund, Charlie Kaufman, Joe Clarke, and Meral Shirazipour performed reviews or provided substantial comments.

Author's Address

Alexander Mayrhofer
nic.at GmbH
Karlsplatz 1/2/9
Vienna 1010
Austria

Email: alex.mayrhofer.ietf@gmail.com
URI: <http://edns0-padding.org/>