DNS Extensions (DNSEXT)
Internet-Draft
Intended status: Standards Track
Expires: October 22, 2009

A. Hubert
Netherlabs Computer Consulting BV.
D. Ulevitch
EveryDNS

April 20, 2009

EDNS Option for performing a data PING
draft-hubert-ulevitch-edns-ping-01.txt

Status of this Memo

Copyright Notice

Abstract

For various reasons, it may be desirable to ask a remote nameserver to add certain data to the response to a query.

This document describes an EDNS option that implements such behavioiur.

Table of Contents

1. Key words

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Introduction

This document describes an EDNS option that can be used to ask a remote nameserver, be it authoritative or a caching resolver, to copy an opaque string from the query to the response.

This string can be used to verify proper transmission of DNS questions and responses of various sizes.

Additionally, implementations could utilise EDNS PING as a way to enhance the security of DNS over UDP.

3. Protocol

This document uses an EDNS [RFC2671] option to signal that the remote nameserver must copy this option, and its payload, from the query to the response, without truncation or modification.

3.1. Nameserver Behaviour

A name server that understands the PING option and chooses to honor a particular PING request MUST respond by including the opaque payload in a PING option in an EDNS OPT pseudo-RR in the response message.

The PING response should be included in addition to the records that would be returned if no PING request were included.

An oversized payload MUST be ignored.

3.2. Resolver Behaviour

Resolvers, including stub resolvers, can signal their desire for an EDNS PING response by adding a PING option in an EDNS OPT pseudo-RR in the question message.

The resolver is free to choose a length for the opaque payload of the PING option request, but care should be taken not to exceed acceptable DNS packet size limits.

Malformed or truncated responses should be treated as suspicicous. Empty responses, however, may simply indicate a response from a nameserver which does not support EDNS PING responses.

3.3. The PING option

The OPTION-CODE for the PING option is 5.

The OPTION-DATA for the PING option is an opaque byte string, the semantics of which are deliberately left outside of this document.

The minimum length of the OPTION-DATA is 4 bytes, the maximum length is 16 bytes.

3.4. Presentation format

The presentation format of the PING option is left outside the scope of the protocol. It should be observed that the payload of the PING option is completely arbitrary, and need not be null-terminated, and in general will not be.

4. Discussion

The PING option is modeled on ICMP ECHO-REQUEST and ECHO-RESPONSE packets ([RFC0792]), and can in fact be used in a similar manner to verify connectivity.

An example of such verification is to determine the maximum response size that arrives unscathed.

In addition, a resolver is free to append a PING option to outgoing queries in order to protect itself from accepting false data by requesting a more clearly marked response. Such a PING-adorned response can clearly be separated from responses sent by third parties.

4.1. Truncation

In some cases, adding the PING option to a response message may trigger message truncation. This specification does not change the rules for DNS message truncation in any way, but implementers will need to pay attention to this issue.

Implementations claiming conformance to this draft, and which are configured to honor PING requests MUST respond to such requests, and must not drop the PING response to prevent truncation.

By definition, a resolver that requests PING responses also supports EDNS, so a resolver that requests PING responses can also use the "sender's UDP payload size" field of the OPT pseudo-RR to signal a receive buffer size large enough to make truncation unlikely.

5. Possible Uses and Implementation Guidance

While this document standardizes how the EDNS PING option can be used, it does not specify how or when it should be used.

In this non-normative section, guidance is given how this option might best be used to achieve certain effects. It is expected that this guidance will be supplanted by the experience of implementors

over time.

In case the EDNS-PING option is used to protect against the spoofing of DNS answers, care must be taken that the payload of the EDNS-PING is sufficiently long and sufficiently unpredictable to serve this purpose.

Proper unpredictability can be achieved by employing a high quality (pseudo-)random generator, as described in [RFC4086].

Not all servers support EDNS Options, nor do all servers respond well to EDNS queries per se. Like EDNS in general, care must be taken to determine if a nameserver responds well to EDNS-PING adorned queries.

If the state of a remote server's support for EDNS-PING is cached, and EDNS-PING is used to protect against spoofing, it is imperative that such state can not be downgraded within a reasonable timeframe.

6. Security Considerations

While EDNS PING might be used to enhance the security of query/ response correlation, in and of itself it is not expected to have security implications.

7. IANA Considerations

IANA is expected and requested to reserve option 5 for EDNS PING.

8. Acknowledgments

Donald Eastlake first discussed the concept of DNS cookies ([I-D.eastlake-dnsext-cookies]), which are remarkably similar to EDNS PING requests, but cover a wider scope and have a defined purpose.

Most of this document was copied almost verbatim from [RFC5001] which implements a very similar EDNS option, used for very different purposes. Thanks are due to Rob Austein and other contributors to the NSID RFC.

Although any mistakes remain our own, the authors gratefully acknowledge the help and contributions of:

Peter van Dijk,

Aki Tuomi

9. References

9.1. Normative References

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, July 1997.

[RFC2671] Vixie, P., "Extension Mechanisms for DNS (EDNS0)", RFC 2671, August 1999.

9.2. Informative References

[I-D.eastlake-dnsext-cookies] 3rd, D., "Domain Name System (DNS) Cookies", draft-eastlake-dnsext-cookies-03 (work in progress), February 2008.

[RFC1123] Braden, R., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, October 1989.

[RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.

[RFC5001] Austein, R., "DNS Name Server Identifier (NSID) Option", RFC 5001, August 2007.

Authors' Addresses

Bert Hubert
Netherlabs Computer Consulting BV.
Braillelaan 10
Rijswijk (ZH) 2289 CM
The Netherlands

Email: bert.hubert@netherlabs.nl

David Ulevitch
EveryDNS
2601 Greenwich, #4
San Francisco, CA 94123
United States of America

Email: davidu@everydns.net